

PCT

WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 6 : G06K 19/14		A2	(11) International Publication Number: WO 95/14289 (43) International Publication Date: 26 May 1995 (26.05.95)
(21) International Application Number: PCT/US94/13366		(74) Agent: CONWELL, William, Y.; Klarquist, Sparkman, Campbell, Leigh & Whinston, One World Trade Center, Suite 1600, 121 S.W. Salmon Street, Portland, OR 97204 (US).	
(22) International Filing Date: 16 November 1994 (16.11.94)		(81) Designated States: CA, JP, US, European patent (AT, BE, CH, DE, DK, ES, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).	
(30) Priority Data: 154,866 18 November 1993 (18.11.93) US 215,289 17 March 1994 (17.03.94) US 327,426 21 October 1994 (21.10.94) US		Published <i>Without international search report and to be republished upon receipt of that report.</i>	
(60) Parent Application or Grant (63) Related by Continuation US 08/327,426 (CIP) Filed on 21 October 1994 (21.10.94)			
(71) Applicant (for all designated States except US): PINECONE IMAGING CORPORATION [US/US]; 363 S.W. Tualatin Loop, West Linn, OR 97068 (US).			
(72) Inventor; and (75) Inventor/Applicant (for US only): RHOADS, Geoffrey, B. [US/US]; 363 S.W. Tualatin Loop, West Linn, OR 97068 (US).			

(54) Title: IDENTIFICATION/AUTHENTICATION CODING METHOD AND APPARATUS

(57) Abstract

An identification code signal is impressed on a carrier to be identified (such as an electronic data signal or a physical medium) in a manner that permits the identification signal later to be discerned and the carrier thereby identified. The method and apparatus are characterized by robustness despite degradation of the encoded carrier, and by holographic permeation of the identification signal throughout the carrier. An exemplary embodiment is a processor that embeds the identification signal onto a carrier signal in real time.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	GB	United Kingdom	MR	Mauritania
AU	Australia	GE	Georgia	MW	Malawi
BB	Barbados	GN	Guinea	NE	Niger
BE	Belgium	GR	Greece	NL	Netherlands
BF	Burkina Faso	HU	Hungary	NO	Norway
BG	Bulgaria	IE	Ireland	NZ	New Zealand
BJ	Benin	IT	Italy	PL	Poland
BR	Brazil	JP	Japan	PT	Portugal
BY	Belarus	KE	Kenya	RO	Romania
CA	Canada	KG	Kyrgyzstan	RU	Russian Federation
CF	Central African Republic	KP	Democratic People's Republic of Korea	SD	Sudan
CG	Congo	KR	Republic of Korea	SE	Sweden
CH	Switzerland	KZ	Kazakhstan	SI	Slovenia
CI	Côte d'Ivoire	LI	Liechtenstein	SK	Slovakia
CM	Cameroon	LK	Sri Lanka	SN	Senegal
CN	China	LU	Luxembourg	TD	Chad
CS	Czechoslovakia	LV	Latvia	TG	Togo
CZ	Czech Republic	MC	Monaco	TJ	Tajikistan
DE	Germany	MD	Republic of Moldova	TT	Trinidad and Tobago
DK	Denmark	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	US	United States of America
FI	Finland	MN	Mongolia	UZ	Uzbekistan
FR	France			VN	Viet Nam
GA	Gabon				

IDENTIFICATION/AUTHENTICATION CODING METHOD AND APPARATUSField of the Invention

The present invention relates to the embedding of robust identification codes in electronic, optical and physical media, and the subsequent, objective discernment of such codes for identification purposes even after intervening distortion or corruption of the media.

The invention is illustrated with reference to several exemplary applications, including identification/authentication coding of electronic imagery, serial data signals (e.g. audio and video), emulsion film, and paper currency, but is not so limited.

Background and Summary of the Invention

10 *"I would never put it in the power of any printer or publisher
to suppress or alter a work of mine, by making him master of
the copy"*

Thomas Paine, *Rights of Man*, 1792.

15 *"The printer dares not go beyond his licensed copy"*
Milton, *Aeropagetica*, 1644.

Since time immemorial, unauthorized use and outright piracy of proprietary source material has been a source of lost revenue, confusion, and artistic corruption.

20 These historical problems have been compounded by the advent of digital technology. With it, the technology of copying materials and redistributing them in unauthorized manners has reached new heights of sophistication, and more importantly, omnipresence. Lacking objective means for comparing an alleged copy of material with the original, owners and possible litigation proceedings are left with a subjective opinion of whether the alleged copy is stolen, or
25 has been used in an unauthorized manner. Furthermore, there is no simple means of tracing a path to an original purchaser of the material, something which can be valuable in tracing where a possible "leak" of the material first occurred.

30 A variety of methods for protecting commercial material have been attempted. One is to scramble signals via an encoding method prior to distribution, and descramble prior to use. This technique, however, requires that both the original and later descrambled signals never leave closed and controlled networks, lest they be intercepted and recorded. Furthermore, this arrangement is of little use in the broad field of mass marketing audio and visual material, where even a few dollars extra cost causes a major reduction in market, and where the signal must eventually be descrambled to be perceived, and thus can be easily recorded.

35 Another class of techniques relies on modification of source audio or video signals to include a subliminal identification signal, which can be sensed by electronic means. Examples of such systems are found in U.S. Patent 4,972,471 and European patent publication EP 441,702, as well as in Komatsu et al, "Authentication System Using Concealed Image in Telematics," Memoirs of the School of Science & Engineering, Waseda University, No. 52, p. 45-
40 60 (1988) (Komatsu uses the term "digital watermark" for this technique). An elementary introduction to these methods is found in the article "Digital Signatures," Byte Magazine,

November, 1993, p. 309. These techniques have the common characteristic that deterministic signals with well defined patterns and sequences within the source material convey the identification information. For certain applications this is not a drawback. But in general, this is an inefficient form of embedding identification information for a variety of reasons: (a) the whole of the source material is not used; (b) deterministic patterns have a higher likelihood of being discovered and removed by a would-be pirate; and (c) the signals are not generally 'holographic' in that identifications may be difficult to make given only sections of the whole. ('Holographic' is used herein to refer to the property that the identification information is distributed globally throughout the coded signal, and can be fully discerned from an examination of even a fraction of the coded signal. Coding of this type is sometimes termed "distributed" herein.)

Among the cited references are descriptions of several programs which perform steganography - described in one document as "... the ancient art of hiding information in some otherwise inconspicuous information." These programs variously allow computer users to hide their own messages inside digital image files and digital audio files. All do so by toggling the least significant bit (the lowest order bit of a single data sample) of a given audio data stream or rasterized image. Some of these programs embed messages quite directly into the least significant bit, while other "pre-encrypt" or scramble a message first and then embed the encrypted data into the least significant bit.

Our current understanding of these programs is that they generally rely on error-free transmission of the of digital data in order to correctly transmit a given message in its entirety. Typically the message is passed only once, i.e., it is not repeated. These programs also seem to "take over" the least significant bit entirely, where actual data is obliterated and the message placed accordingly. This might mean that such codes could be easily erased by merely stripping off the least significant bit of all data values in a given image or audio file. It is these and other considerations which suggest that the only similarity between our invention and the established art of steganography is in the placement of information into data files with minimal perceptibility. The specifics of embedding and the uses of that buried information diverge from there.

Another cited reference is U.S. Patent 5,325,167 to Melen. In the service of authenticating a given document, the high precision scanning of that document reveals patterns and "microscopic grain structure" which apparently is a kind of unique fingerprint for the underlying document media, such as paper itself or post-applied materials such as toner. Melen further teaches that scanning and storing this fingerprint can later be used in authentication by scanning a purported document and comparing it to the original fingerprint. Applicant is aware of a similar idea employed in the very high precision recording of credit card magnetic strips, as reported in the February 8, 1994, Wall Street Journal, page B1, wherein very fine magnetic fluxuations tend to be unique from one card to the next, so that credit card authentication could be achieved

through pre-recording these fluxuations later to be compared to the recordings of the purportedly same credit card.

Both of the foregoing techniques appear to rest on the same identification principles on which the mature science of fingerprint analysis rests: the innate uniqueness of some 5 localized physical property. These methods then rely upon a single judgement and/or measurement of "similarity" or "correlation" between a suspect and a pre-recording master. Though fingerprint analysis has brought this to a high art, these methods are nevertheless open to a claim that preparations of the samples, and the "filtering" and "scanner specifications" of Melen's patent, unavoidably tend to bias the resulting judgement of similarity, and would create a 10 need for more esoteric "expert testimony" to explain the confidence of a found match or mis-match. An object of the present invention is to avoid this reliance on expert testimony and to place the confidence in a match into simple "coin flip" vernacular, i.e., what are the odds you can call the correct coin flip 16 times in a row. Attempts to identify fragments of a fingerprint, document, or otherwise, exacerbate this issue of confidence in a judgment, where it is an object of 15 the present invention to objectively apply the intuitive "coin flip" confidence to the smallest fragment possible. Also, storing unique fingerprints for each and every document or credit card magnetic strip, and having these fingerprints readily available for later cross-checking, should prove to be quite an economic undertaking. It is an object of this invention to allow for the "re-use" of noise codes and "snowy images" in the service of easing storage requirements.

20 U.S. Patent 4,921,278 to Shiang et al. teaches a kind of spatial encryption technique wherein a signature or photograph is splayed out into what the untrained eye would refer to as noise, but which is actually a well defined structure referred to as Moire patterns. The similarities of the present invention to Shiang's system appear to be use of noise-like patterns which nevertheless carry information, and the use of this principle on credit cards and other 25 identification cards.

Others of the cited patents deal with other techniques for identification and/or authentication of signals or media. U.S. Patent 4,944,036 to Hyatt does not appear to be applicable to the present invention, but does point out that the term "signature" can be equally applied to signals which carry unique characteristics based on physical structure.

30 Despite the foregoing and other diverse work in the field of identification/authentication, there still remains a need for a reliable and efficient method for performing a positive identification between a copy of an original signal and the original. Desirably, this method should not only perform identification, it should also be able to convey source-version information in order to better pinpoint the point of sale. The method should not 35 compromise the innate quality of material which is being sold, as does the placement of localized logos on images. The method should be robust so that an identification can be made even after multiple copies have been made and/or compression and decompression of the signal has taken place. The identification method should be largely uneraseable or "uncrackable." The method

should be capable of working even on fractional pieces of the original signal, such as a 10 second "tiff" of an audio signal or the "clipped and pasted" sub-section of an original image.

The existence of such a method would have profound consequences on piracy in that it could (a) cost effectively monitor for unauthorized uses of material and perform "quick checks"; (b) become a deterrent to unauthorized uses when the method is known to be in use and the consequences well publicized; and (c) provide unequivocal proof of identity, similar to fingerprint identification, in litigation, with potentially more reliability than that of fingerprinting.

In accordance with an exemplary embodiment of the invention, the foregoing and additional objects are achieved by embedding an imperceptible identification code throughout a source signal. In the preferred embodiment, this embedding is achieved by modulating the source signal with a small noise signal in a coded fashion. More particularly, bits of a binary identification code are referenced, one at a time, to control modulation of the source signal with the noise signal.

The copy with the embedded signal (the "encoded" copy) becomes the material which is sold, while the original is secured in a safe place. The new copy is nearly identical to the original except under the finest of scrutiny; thus, its commercial value is not compromised. After the new copy has been sold and distributed and potentially distorted by multiple copies, the present disclosure details methods for positively identifying any suspect signal against the original.

Among its other advantages, the preferred embodiments' use of identification signals which are global (holographic) and which mimic natural noise sources allows the maximization of identification signal energy, as opposed to merely having it present 'somewhere in the original material.' This allows the identification coding to be much more robust in the face of thousands of real world degradation processes and material transformations, such as cutting and cropping of imagery.

The foregoing and additional features and advantages of the present invention will be more readily apparent from the following detailed description thereof, which proceeds with reference to the accompanying drawings.

Brief Description of the Drawings

Fig. 1 is a simple and classic depiction of a one dimensional digital signal which is discretized in both axes.

Fig. 2 is a general overview, with detailed description of steps, of the process of embedding an "imperceptible" identification signal onto another signal.

Fig. 3 is a step-wise description of how a suspected copy of an original is identified.

Fig. 4 is a schematic view of an apparatus for pre-exposing film with identification information in accordance with another embodiment of the present invention.

Fig. 5 is a diagram of a "black box" embodiment of the present invention.

Fig. 6 is a schematic block diagram of the embodiment of Fig. 5.

Fig. 7 shows a variant of the Fig. 6 embodiment adapted to encode successive sets of input data with different code words but with the same noise data.

Fig. 8 shows a variant of the Fig. 6 embodiment adapted to encode each frame of a videotaped production with a unique code number.

5 Figs. 9A-9C are representations of an industry standard noise second that can be used in one embodiment of the present invention.

Fig. 10 shows an integrated circuit used in detecting standard noise codes.

Fig. 11 shows a process flow for detecting a standard noise code that can be used in the Fig. 10 embodiment.

10 Fig. 12 is an embodiment employing a plurality of detectors in accordance with another embodiment of the present invention.

Detailed Description

In the following discussion of an illustrative embodiment, the words "signal" and "image" are used interchangeably to refer to both one, two, and even beyond two dimensions of 15 digital signal. Examples will routinely switch back and forth between a one dimensional audio-type digital signal and a two dimensional image-type digital signal.

In order to fully describe the details of an illustrative embodiment of the invention, it is necessary first to describe the basic properties of a digital signal. Fig. 1 shows a classic representation of a one dimensional digital signal. The x-axis defines the index numbers of 20 sequence of digital "samples," and the y-axis is the instantaneous value of the signal at that sample, being constrained to exist only at a finite number of levels defined as the "binary depth" of a digital sample. The example depicted in Fig. 1 has the value of 2 to the fourth power, or "4 bits," giving 16 allowed states of the sample value.

For audio information such as sound waves, it is commonly accepted that the 25 digitization process discretizes a continuous phenomena both in the time domain and in the signal level domain. As such, the process of digitization itself introduces a fundamental error source, in that it cannot record detail smaller than the discretization interval in either domain. The industry has referred to this, among other ways, as "aliasing" in the time domain, and "quantization noise" in the signal level domain. Thus, there will always be a basic error floor of a digital signal. Pure 30 quantization noise, measured in a root mean square sense, is theoretically known to have the value of one over the square root of twelve, or about 0.29 DN, where DN stands for 'Digital Number' or the finest unit increment of the signal level. For example, a perfect 12-bit digitizer will have 4096 allowed DN with an innate root mean square noise floor of ~0.29 DN.

All known physical measurement processes add additional noise to the 35 transformation of a continuous signal into the digital form. The quantization noise typically adds in quadrature (square root of the mean squares) to the "analog noise" of the measurement process, as it is sometimes referred to.

With almost all commercial and technical processes, the use of the decibel scale is used as a measure of signal and noise in a given recording medium. The expression "signal-to-noise ratio" is generally used, as it will be in this disclosure. As an example, this disclosure refers to signal to noise ratios in terms of signal power and noise power, thus 20 dB represents a 10 times increase in signal amplitude.

- 5 In summary, the presently preferred embodiments of the invention embed an N-bit value onto an entire signal through the addition of a very low amplitude encodation signal which has the look of pure noise. N is usually at least 8 and is capped on the higher end by ultimate signal-to-noise considerations and "bit error" in retrieving and decoding the N-bit value.
- 10 10 As a practical matter, N is chosen based on application specific considerations, such as the number of unique different "signatures" that are desired. To illustrate, if N=128, then the number of unique digital signatures is in excess of 10^{38} (2^{128}). This number is believed to be more than adequate to both identify the material with sufficient statistical certainty and to index exact sale and distribution information.
- 15 15 The amplitude or power of this added signal is determined by the aesthetic and informational considerations of each and every application using the present methodology. For instance, non-professional video can stand to have a higher embedded signal level without becoming noticeable to the average human eye, while high precision audio may only be able to accept a relatively small signal level lest the human ear perceive an objectionable increase in "hiss." These statements are generalities and each application has its own set of criteria in choosing the signal level of the embedded identification signal. The higher the level of embedded signal, the more corrupted a copy can be and still be identified. On the other hand, the higher the level of embedded signal, the more objectionable the perceived noise might be, potentially impacting the value of the distributed material.
- 20 20 25 To illustrate the range of different applications to which the principles of the present invention can be applied, the present specification details two different systems. The first (termed, for lack of a better name, a "batch encoding" system), applies identification coding to an existing data signal. The second (termed, for lack of a better name, a "real time encoding" system), applies identification coding to a signal as it is produced. Those skilled in the art will recognize that the principles of the present invention can be applied in a number of other contexts in addition to these particularly described.

The discussions of these two systems can be read in either order. Some readers may find the latter more intuitive than the former; for others the contrary may be true.

BATCH ENCODING

- 30 35 The following discussion of a first class of embodiments is best prefaced by a section defining relevant terms:

The original signal refers to either the original digital signal or the high quality digitized copy of a non-digital original.

5 The N-bit identification word refers to a unique identification binary value, typically having N range anywhere from 8 to 128, which is the identification code ultimately placed onto the original signal via the disclosed transformation process. In the illustrated embodiment, each N-bit identification word begins with the sequence of values '0101,' which is used to determine an optimization of the signal-to-noise ratio in the identification procedure of a suspect signal (see definition below).

10 The m'th bit value of the N-bit identification word is either a zero or one corresponding to the value of the m'th place, reading left to right, of the N-bit word. E.g., the first (m=1) bit value of the N=8 identification word 01110100 is the value '0'; the second bit value of this identification word is '1', etc.

15 The m'th individual embedded code signal refers to a signal which has dimensions and extent precisely equal to the original signal (e.g. both are a 512 by 512 digital image), and which is (in the illustrated embodiment) an independent pseudo-random sequence of digital values. "Pseudo" pays homage to the difficulty in philosophically defining pure randomness, and also indicates that there are various acceptable ways of generating the "random" signal. There will be exactly N individual embedded code signals associated with any given original signal.

20 The acceptable perceived noise level refers to an application-specific determination of how much "extra noise," i.e. amplitude of the composite embedded code signal described next, can be added to the original signal and still have an acceptable signal to sell or otherwise distribute. This disclosure uses a 1 dB increase in noise as a typical value which might be acceptable, but this is quite arbitrary.

25 The composite embedded code signal refers to the signal which has dimensions and extent precisely equal to the original signal, (e.g. both are a 512 by 512 digital image), and which contains the addition and appropriate attenuation of the N individual embedded code signals. The individual embedded signals are generated on an arbitrary scale, whereas the amplitude of the composite signal must not exceed the pre-set acceptable perceived noise level, hence the need for "attenuation" of the N added individual code signals.

30 The distributable signal refers to the nearly similar copy of the original signal, consisting of the original signal plus the composite embedded code signal. This is the signal which is distributed to the outside community, having only slightly higher but acceptable "noise properties" than the original.

35 A suspect signal refers to a signal which has the general appearance of the original and distributed signal and whose potential identification match to the original is being questioned. The suspect signal is then analyzed to see if it matches the N-bit identification word.

 The detailed methodology of this first embodiment begins by stating that the N-bit identification word is encoded onto the original signal by having each of the m bit values multiply their corresponding individual embedded code signals, the resultant being accumulated in

the composite signal, the fully summed composite signal then being attenuated down to the acceptable perceived noise amplitude, and the resultant composite signal added to the original to become the distributable signal.

5 The original signal, the N-bit identification word, and all N individual embedded code signals are then stored away in a secured place. A suspect signal is then found. This signal may have undergone multiple copies, compressions and decompressions, resamplings onto different spaced digital signals, transfers from digital to analog back to digital media, or any combination of these items. IF the signal still appears similar to the original, i.e. its innate quality is not thoroughly destroyed by all of these transformations and noise additions, then depending on the
10 signal to noise properties of the embedded signal, the identification process should function to some objective degree of statistical confidence. The extent of corruption of the suspect signal and the original acceptable perceived noise level are two key parameters in determining an expected confidence level of identification.

15 The identification process on the suspected signal begins by resampling and aligning the suspected signal onto the digital format and extent of the original signal. Thus, if an image has been reduced by a factor of two, it needs to be digitally enlarged by that same factor. Likewise, if a piece of music has been "cut out," but may still have the same sampling rate as the original, it is necessary to register this cut-out piece to the original, typically done by performing a local digital cross-correlation of the two signals (a common digital operation), finding at what
20 delay value the correlation peaks, then using this found delay value to register the cut piece to a segment of the original.

Once the suspect signal has been sample-spacing matched and registered to the original, the signal levels of the suspect signal should be matched in an rms sense to the signal level of the original. This can be done via a search on the parameters of offset, amplification, and
25 gamma being optimized by using the minimum of the mean squared error between the two signals as a function of the three parameters. We can call the suspect signal normalized and registered at this point, or just normalized for convenience.

The newly matched pair then has the original signal subtracted from the normalized suspect signal to produce a difference signal. The difference signal is then
30 cross-correlated with each of the N individual embedded code signals and the peak cross-correlation value recorded. The first four bit code ('0101') is used as a calibrator both on the mean values of the zero value and the one value, and on further registration of the two signals if a finer signal to noise ratio is desired (i.e., the optimal separation of the 0101 signal will indicate an optimal registration of the two signals and will also indicate the probable existence of
35 the N-bit identification signal being present.)

The resulting peak cross-correlation values will form a noisy series of floating point numbers which can be transformed into 0's and 1's by their proximity to the mean values of 0 and 1 found by the 0101 calibration sequence. If the suspect signal has indeed been derived

-9-

from the original, the identification number resulting from the above process will match the N-bit identification word of the original, bearing in mind either predicted or unknown "bit error" statistics. Signal-to-noise considerations will determine if there will be some kind of "bit error" in the identification process, leading to a form of X% probability of identification where X might be
5 desired to be 99.9% or whatever. If the suspect copy is indeed not a copy of the original, an essentially random sequence of 0's and 1's will be produced, as well as an apparent lack of separation of the resultant values. This is to say, if the resultant values are plotted on a histogram, the existence of the N-bit identification signal will exhibit strong bi-level characteristics, whereas the non-existence of the code, or the existence of a different code of a different original, will
10 exhibit a type of random gaussian-like distribution. This histogram separation alone should be sufficient for an identification, but it is even stronger proof of identification when an exact binary sequence can be objectively reproduced.

Specific Example

Imagine that we have taken a valuable picture of two heads of state at a cocktail party, pictures which are sure to earn some reasonable fee in the commercial market. We desire
15 to sell this picture and ensure that it is not used in an unauthorized or uncompensated manner. This and the following steps are summarized in Fig. 2.

Assume the picture is transformed into a positive color print. We first scan this into a digitized form via a normal high quality black and white scanner with a typical photometric
20 spectral response curve. (It is possible to get better ultimate signal to noise ratios by scanning in each of the three primary colors of the color image, but this nuance is not central to describing the basic process.)

Let us assume that the scanned image now becomes a 4000 by 4000 pixel monochrome digital image with a grey scale accuracy defined by 12-bit grey values or 4096
25 allowed levels. We will call this the "original digital image" realizing that this is the same as our "original signal" in the above definitions.

During the scanning process we have arbitrarily set absolute black to correspond to digital value '30'. We estimate that there is a basic 2 Digital Number root mean square noise existing on the original digital image, plus a theoretical noise (known in the industry as "shot
30 noise") of the square root of the brightness value of any given pixel. In formula, we have:

$$\langle \text{RMS Noise}_{n,m} \rangle = \sqrt{4 + (V_{n,m} - 30)} \quad (1)$$

Here, n and m are simple indexing values on rows and columns of the image ranging from 0 to
35 3999. Sqrt is the square root. V is the DN of a given indexed pixel on the original digital image. The < > brackets around the RMS noise merely indicates that this is an expected average value, where it is clear that each and every pixel will have a random error individually. Thus, for a pixel

-10-

value having 1200 as a digital number or "brightness value", we find that its expected rms noise value is $\sqrt{1204} = 34.70$, which is quite close to 34.64, the square root of 1200.

5 We furthermore realize that the square root of the innate brightness value of a pixel is not precisely what the eye perceives as a minimum objectionable noise, thus we come up with the formula:

$$\langle \text{RMS Addable Noise}_{n,m} \rangle = X * \sqrt{4 + (V_{n,m} - 30)^Y} \quad (2)$$

10 Where X and Y have been added as empirical parameters which we will adjust, and "addable" noise refers to our acceptable perceived noise level from the definitions above. We now intend to experiment with what exact value of X and Y we can choose, but we will do so at the same time that we are performing the next steps in the process.

15 The next step in our process is to choose N of our N-bit identification word. We decide that a 16 bit main identification value with its 65536 possible values will be sufficiently large to identify the image as ours, and that we will be directly selling no more than 128 copies of the image which we wish to track, giving 7 bits plus an eighth bit for an odd/even adding of the first 7 bits (i.e. an error checking bit on the first seven). The total bits required now are at 4 bits for the 0101 calibration sequence, 16 for the main identification, 8 for the version, and we now throw in another 4 as a further error checking value on the first 28 bits, giving 32 bits as N. The 20 final 4 bits can use one of many industry standard error checking methods to choose its four values.

25 We now randomly determine the 16 bit main identification number, finding for example, 1101 0001 1001 1110; our first versions of the original sold will have all 0's as the version identifier, and the error checking bits will fall out where they may. We now have our unique 32 bit identification word which we will embed on the original digital image.

30 To do this, we generate 32 independent random 4000 by 4000 encoding images for each bit of our 32 bit identification word. The manner of generating these random images is revealing. There are numerous ways to generate these. By far the simplest is to turn up the gain on the same scanner that was used to scan in the original photograph, only this time placing a pure black image as the input, then scanning this 32 times. The only drawback to this technique is that it does require a large amount of memory and that "fixed pattern" noise will be part of each independent "noise image." But, the fixed pattern noise can be removed via normal "dark frame" subtraction techniques. Assume that we set the absolute black average value at digital number '100,' and that rather than finding a 2 DN rms noise as we did in the normal gain setting, we now 35 find an rms noise of 10 DN about each and every pixel's mean value.

We next apply a mid-spatial-frequency bandpass filter (spatial convolution) to each and every independent random image, essentially removing the very high and the very low spatial frequencies from them. We remove the very low frequencies because simple real-world

-11-

error sources like geometrical warping, splotches on scanners, mis-registrations, and the like will exhibit themselves most at lower frequencies also, and so we want to concentrate our identification signal at higher spatial frequencies in order to avoid these types of corruptions. Likewise, we remove the higher frequencies because multiple generation copies of a given image, as well as compression-decompression transformations, tend to wipe out higher frequencies anyway, so there is no point in placing too much identification signal into these frequencies if they will be the ones most prone to being attenuated. Therefore, our new filtered independent noise images will be dominated by mid-spatial frequencies. On a practical note, since we are using 12-bit values on our scanner and we have removed the DC value effectively and our new rms noise will be slightly less than 10 digital numbers, it is useful to boil this down to a 6-bit value ranging from -32 through 0 to 31 as the resultant random image.

10 Next we add all of the random images together which have a '1' in their corresponding bit value of the 32-bit identification word, accumulating the result in a 16-bit signed integer image. This is the unattenuated and un-scaled version of the composite embedded signal.

15 Next we experiment visually with adding the composite embedded signal to the original digital image, through varying the X and Y parameters of equation 2. In formula, we visually iterate to both maximize X and to find the appropriate Y in the following:

$$V_{dist,n,m} = V_{orig,n,m} + V_{comp,n,m} * X * \sqrt{4 + V_{orig,n,m}^2} \quad (3)$$

20 where dist refers to the candidate distributable image, i.e. we are visually iterating to find what X and Y will give us an acceptable image; orig refers to the pixel value of the original image; and comp refers to the pixel value of the composite image. The n's and m's still index rows and columns of the image and indicate that this operation is done on all 4000 by 4000 pixels. The symbol V is the DN of a given pixel and a given image.

25 As an arbitrary assumption, now, we assume that our visual experimentation has found that the value of X= 0.025 and Y=0.6 are acceptable values when comparing the original image with the candidate distributable image. This is to say, the distributable image with the "extra noise" is acceptably close to the original in an aesthetic sense. Note that since our individual random images had a random rms noise value around 10 DN, and that adding approximately 16 of these images together will increase the composite noise to around 40 DN, the X multiplication value of 0.025 will bring the added rms noise back to around 1 DN, or half the amplitude of our innate noise on the original. This is roughly a 1 dB gain in noise at the dark pixel values and correspondingly more at the brighter values modified by the Y value of 0.6.

30 So with these two values of X and Y, we now have constructed our first versions of a distributable copy of the original. Other versions will merely create a new composite signal and possibly change the X slightly if deemed necessary. We now lock up the original digital image along with the 32-bit identification word for each version, and the 32 independent random

-12-

4-bit images, waiting for our first case of a suspected piracy of our original. Storage wise, this is about 14 Megabytes for the original image and $32 * 0.5 \text{ bytes} * 16 \text{ million} = \sim 256 \text{ Megabytes}$ for the random individual encoded images. This is quite acceptable for a single valuable image. Some storage economy can be gained by simple lossless compression.

5 Finding a Suspected Piracy of our Image

We sell our image and several months later find our two heads of state in the exact poses we sold them in, seemingly cut and lifted out of our image and placed into another stylized background scene. This new "suspect" image is being printed in 100,000 copies of a given magazine issue, let us say. We now go about determining if a portion of our original image
10 has indeed been used in an unauthorized manner. Fig. 3 summarizes the details.

The first step is to take an issue of the magazine, cut out the page with the image on it, then carefully but not too carefully cut out the two figures from the background image using ordinary scissors. If possible, we will cut out only one connected piece rather than the two figures separately. We paste this onto a black background and scan this into a digital
15 form. Next we electronically flag or mask out the black background, which is easy to do by visual inspection.

We now procure the original digital image from our secured place along with the 32-bit identification word and the 32 individual embedded images. We place the original digital image onto our computer screen using standard image manipulation software, and we roughly cut
20 along the same borders as our masked area of the suspect image, masking this image at the same time in roughly the same manner. The word 'roughly' is used since an exact cutting is not needed, it merely aids the identification statistics to get it reasonably close.

Next we rescale the masked suspect image to roughly match the size of our masked original digital image, that is, we digitally scale up or down the suspect image and
25 roughly overlay it on the original image. Once we have performed this rough registration, we then throw the two images into an automated scaling and registration program. The program performs a search on the three parameters of x position, y position, and spatial scale, with the figure of merit being the mean squared error between the two images given any given scale variable and x and y offset. This is a fairly standard image processing methodology. Typically
30 this would be done using generally smooth interpolation techniques and done to sub-pixel accuracy. The search method can be one of many, where the simplex method is a typical one.

Once the optimal scaling and x-y position variables are found, next comes another search on optimizing the black level, brightness gain, and gamma of the two images. Again, the figure of merit to be used is mean squared error, and again the simplex or other search
35 methodologies can be used to optimize the three variables. After these three variables are optimized, we apply their corrections to the suspect image and align it to exactly the pixel spacing and masking of the original digital image and its mask. We can now call this the standard mask.

-13-

The next step is to subtract the original digital image from the newly normalized suspect image only within the standard mask region. This new image is called the difference image.

Then we step through all 32 individual random embedded images, doing a local cross-correlation between the masked difference image and the masked individual embedded image. 'Local' refers to the idea that one need only start correlating over an offset region of +/- 1 pixels of offset between the nominal registration points of the two images found during the search procedures above. The peak correlation should be very close to the nominal registration point of 0,0 offset, and we can add the 3 by 3 correlation values together to give one grand correlation value for each of the 32 individual bits of our 32-bit identification word.

After doing this for all 32 bit places and their corresponding random images, we have a quasi-floating point sequence of 32 values. The first four values represent our calibration signal of 0101. We now take the mean of the first and third floating point value and call this floating point value '0,' and we take the mean of the second and the fourth value and call this floating point value '1.' We then step through all remaining 28 bit values and assign either a '0' or a '1' based simply on which mean value they are closer to. Stated simply, if the suspect image is indeed a copy of our original, the embedded 32-bit resulting code should match that of our records, and if it is not a copy, we should get general randomness. The third and the fourth possibilities of 3) Is a copy but doesn't match identification number and 4) isn't a copy but does match are, in the case of 3), possible if the signal to noise ratio of the process has plummeted, i.e. the 'suspect image' is truly a very poor copy of the original, and in the case of 4) is basically one chance in four billion since we were using a 32-bit identification number. If we are truly worried about 4), we can just have a second independent lab perform their own tests on a different issue of the same magazine. Finally, checking the error-check bits against what the values give is one final and possibly overkill check on the whole process. In situations where signal to noise is a possible problem, these error checking bits might be eliminated without too much harm.

Benefits

Now that a full description of the first embodiment has been described via a detailed example, it is appropriate to point out the rationale of some of the process steps and their benefits.

The ultimate benefits of the foregoing process are that obtaining an identification number is fully independent of the manners and methods of preparing the difference image. That is to say, the manners of preparing the difference image, such as cutting, registering, scaling, etcetera, cannot increase the odds of finding an identification number when none exists; it only helps the signal-to-noise ratio of the identification process when a true identification number is present. Methods of preparing images for identification can be different from each other even, providing the possibility for multiple independent methodologies for making a match.

-14-

The ability to obtain a match even on sub-sets of the original signal or image is a key point in today's information-rich world. Cutting and pasting both images and sound clips is becoming more common, allowing such an embodiment to be used in detecting a copy even when original material has been thus corrupted. Finally, the signal to noise ratio of matching should 5 begin to become difficult only when the copy material itself has been significantly altered either by noise or by significant distortion; both of these also will affect that copy's commercial value, so that trying to thwart the system can only be done at the expense of a huge decrease in commercial value.

An early conception of this invention was the case where only a single "snowy" 10 image" or random signal was added to an original image, i.e. the case where $N=1$. "Decoding" this signal would involve a subsequent mathematical analysis using (generally statistical) algorithms to make a judgment on the presence or absence of this signal. The reason this approach was abandoned as the preferred embodiment was that there was an inherent gray area in the certainty of detecting the presence or absence of the signal. By moving onward to a multitude 15 of bit planes, i.e. $N > 1$, combined with simple pre-defined algorithms prescribing the manner of choosing between a "0" and a "1", the invention moved the certainty question from the realm of expert statistical analysis into the realm of guessing a random binary event such as a coin flip. This is seen as a powerful feature relative to the intuitive acceptance of this invention in both the courtroom and the marketplace. The analogy which summarizes the inventor's thoughts on this 20 whole question is as follows: The search for a single identification signal amounts to calling a coin flip only once, and relying on arcane experts to make the call; whereas the $N>1$ preferred embodiment of this invention relies on the broadly intuitive principle of correctly calling a coin flip N times in a row. This situation is greatly exacerbated, i.e. the problems of "interpretation" of the presence of a single signal, when images and sound clips get smaller and smaller in extent.

25 Another important reason that the $N>1$ case is the preferred embodiment over the $N=1$ embodiment is that in the $N=1$ case, the manner in which a suspect image is prepared and manipulated has a direct bearing on the likelihood of making a positive identification. Thus, the manner with which an expert makes an identification determination becomes an integral part of that determination. The existence of a multitude of mathematical and statistical approaches to 30 making this determination leave open the possibility that some tests might make positive identifications while others might make negative determinations, inviting further arcane debate about the relative merits of the various identification approaches. The $N>1$ preferred embodiment of this invention avoids this further gray area by presenting a method where no amount of pre-processing of a signal - other than pre-processing which surreptitiously uses knowledge of the 35 private code signals - can increase the likelihood of "calling the coin flip N times in a row."

The fullest expression of the present system will come when it becomes an industry standard and numerous independent groups set up with their own means or 'in-house' brand of applying embedded identification numbers and in their decipherment. Numerous

-15-

independent group identification will further enhance the ultimate objectivity of the method, thereby enhancing its appeal as an industry standard.

Use of True Polarity in Creating the Composite Embedded Code Signal

The foregoing discussion made use of the 0 and 1 formalism of binary technology to accomplish its ends. Specifically, the 0's and 1's of the N-bit identification word directly multiplied their corresponding individual embedded code signal to form the composite embedded code signal (step 8, figure 2). This approach certainly has its conceptual simplicity, but the multiplication of an embedded code signal by 0 along with the storage of that embedded code contains a kind of inefficiency.

It is preferred to maintain the formalism of the 0 and 1 nature of the N-bit identification word, but to have the 0's of the word induce a subtraction of their corresponding embedded code signal. Thus, in step 8 of figure 2, rather than only 'adding' the individual embedded code signals which correspond to a '1' in the N-bit identification word, we will also 'subtract' the individual embedded code signals which correspond to a '0' in the N-bit identification word.

At first glance this seems to add more apparent noise to the final composite signal. But it also increases the energy-wise separation of the 0's from the 1's, and thus the 'gain' which is applied in step 10, figure 2 can be correspondingly lower.

We can refer to this improvement as the use of true polarity. The main advantage of this improvement can largely be summarized as 'informational efficiency.'

'Perceptual Orthogonality' of the Individual Embedded Code Signals

The foregoing discussion contemplates the use of generally random noise-like signals as the individual embedded code signals. This is perhaps the simplest form of signal to generate. However, there is a form of informational optimization which can be applied to the set of the individual embedded signals, which the applicant describes under the rubric 'perceptual orthogonality.' This term is loosely based on the mathematical concept of the orthogonality of vectors, with the current additional requirement that this orthogonality should maximize the signal energy of the identification information while maintaining it below some perceptibility threshold. Put another way, the embedded code signals need not necessarily be random in nature.

Use and Improvements of the First Embodiment in the Field of Emulsion-Based Photography

The foregoing discussion outlined techniques that are applicable to photographic materials. The following section explores the details of this area further and discloses certain improvements which lend themselves to a broad range of applications.

The first area to be discussed involves the pre-application or pre-exposing of a serial number onto traditional photographic products, such as negative film, print paper, transparencies, etc. In general, this is a way to embed *a priori* unique serial numbers (and by implication, ownership and tracking information) into photographic material. The serial numbers themselves would be a permanent part of the normally exposed picture, as opposed to being

relegated to the margins or stamped on the back of a printed photograph, which all require separate locations and separate methods of copying. The 'serial number' as it is called here is generally synonymous with the N-bit identification word, only now we are using a more common industrial terminology.

5 In Figure 2, step 11, the disclosure calls for the storage of the "original [image]" along with code images. Then in figure 3, step 9, it directs that the original be subtracted from the suspect image, thereby leaving the possible identification codes plus whatever noise and corruption has accumulated. Therefore, the previous disclosure made the tacit assumption that there exists an original without the composite embedded signals.

10 Now in the case of selling print paper and other duplication film products, this will still be the case, i.e., an "original" without the embedded codes will indeed exist and the basic methodology of the first embodiment can be employed. The original film serves perfectly well as an 'unencoded original.'

15 However, in the case where pre-exposed negative film is used, the composite embedded signal pre-exists on the original film and thus there will never be an "original" separate from the pre-embedded signal. It is this latter case, therefore, which will be examined a bit more closely, along with observations on how to best use the principles discussed above (the former cases adhering to the previously outlined methods).

20 The clearest point of departure for the case of pre-numbered negative film, i.e. negative film which has had each and every frame pre-exposed with a very faint and unique composite embedded signal, comes at step 9 of figure 3 as previously noted. There are certainly other differences as well, but they are mostly logistical in nature, such as how and when to embed the signals on the film, how to store the code numbers and serial number, etc. Obviously the pre-exposing of film would involve a major change to the general mass production process of creating and packaging film.

25 Fig. 4 has a schematic outlining one potential post-hoc mechanism for pre-exposing film. 'Post-hoc' refers to applying a process after the full common manufacturing process of film has already taken place. Eventually, economies of scale may dictate placing this pre-exposing process directly into the chain of manufacturing film. Depicted in Fig. 4 is what is commonly known as a film writing system. The computer, 106, displays the composite signal produced in step 8, figure 2, on its phosphor screen. A given frame of film is then exposed by imaging this phosphor screen, where the exposure level is generally very faint, i.e. generally imperceptible. Clearly, the marketplace will set its own demands on how faint this should be, that is, the level of added 'graininess' as practitioners would put it. Each frame of film is sequentially exposed, where in general the composite image displayed on the CRT 102 is changed for each and every frame, thereby giving each frame of film a different serial number. The transfer lens 104 highlights the focal conjugate planes of a film frame and the CRT face.

-17-

Getting back to the applying the principles of the foregoing embodiment in the case of pre-exposed negative film... At step 9, figure 3, if we were to subtract the "original" with its embedded code, we would obviously be "erasing" the code as well since the code is an integral part of the original. Fortunately, remedies do exist and identifications can still be made.

5 However, it will be a challenge to artisans who refine this embodiment to have the signal to noise ratio of the identification process in the pre-exposed negative case approach the signal to noise ratio of the case where the un-encoded original exists.

A succinct definition of the problem is in order at this point. Given a suspect picture (signal), find the embedded identification code IF a code exists at al. The problem reduces 10 to one of finding the amplitude of each and every individual embedded code signal within the suspect picture, not only within the context of noise and corruption as was previously explained, but now also within the context of the coupling between a captured image and the codes.

'Coupling' here refers to the idea that the captured image "randomly biases" the cross-correlation.

So, bearing in mind this additional item of signal coupling, the identification 15 process now estimates the signal amplitude of each and every individual embedded code signal (as opposed to taking the cross-correlation result of step 12, figure 3). If our identification signal exists in the suspect picture, the amplitudes thus found will split into a polarity with positive amplitudes being assigned a '1' and negative amplitudes being assigned a '0'. Our unique identification code manifests itself. If, on the other hand, no such identification code exists or it is 20 someone else's code, then a random gaussian-like distribution of amplitudes is found with a random hash of values.

It remains to provide a few more details on how the amplitudes of the individual 25 embedded codes are found. Again, fortunately, this exact problem has been treated in other technological applications. Besides, throw this problem and a little food into a crowded room of mathematicians and statisticians and surely a half dozen optimized methodologies will pop out after some reasonable period of time. It is a rather cleanly defined problem.

One specific example solution comes from the field of astronomical imaging. Here, it is a mature prior art to subtract out a "thermal noise frame" from a given CCD image of 30 an object. Often, however, it is not precisely known what scaling factor to use in subtracting the thermal frame, and a search for the correct scaling factor is performed. This is precisely the task of this step of the present embodiment.

General practice merely performs a common search algorithm on the scaling factor, where a scaling factor is chosen and a new image is created according to:

$$\text{NEW IMAGE} = \text{ACQUIRED IMAGE} - \text{SCALE} * \text{THERMAL IMAGE} \quad (4)$$

35 The new image is applied to the fast fourier transform routine and a scale factor is eventually found which minimizes the integrated high frequency content of the new image. This general type of search operation with its minimization of a particular quantity is exceedingly common. The scale factor thus found is the sought-for "amplitude." Refinements which are

contemplated but not yet implemented are where the coupling of the higher derivatives of the acquired image and the embedded codes are estimated and removed from the calculated scale factor. In other words, certain bias effects from the coupling mentioned earlier are present and should be eventually accounted for and removed both through theoretical and empirical
5 experimentation.

Use and Improvements in the Detection of Signal or Image Alteration

Apart from the basic need of identifying a signal or image as a whole, there is also a rather ubiquitous need to detect possible alterations to a signal or image. The following section describes how the foregoing embodiment, with certain modifications and improvements,
10 can be used as a powerful tool in this area. The potential scenarios and applications of detecting alterations are innumerable.

To first summarize, assume that we have a given signal or image which has been positively identified using the basic methods outlined above. In other words, we know its N-bit identification word, its individual embedded code signals, and its composite embedded code. We
15 can then fairly simply create a spatial map of the composite code's amplitude within our given signal or image. Furthermore, we can divide this amplitude map by the known composite code's spatial amplitude, giving a normalized map, i.e. a map which should fluctuate about some global mean value. By simple examination of this map, we can visually detect any areas which have
20 been significantly altered wherein the value of the normalized amplitude dips below some statistically set threshold based purely on typical noise and corruption (error).

The details of implementing the creation of the amplitude map have a variety of choices. One is to perform the same procedure which is used to determine the signal amplitude as described above, only now we step and repeat the multiplication of any given area of the signal/image with a gaussian weight function centered about the area we are investigating.
25

Universal Versus Custom Codes

The disclosure thus far has outlined how each and every source signal has its own unique set of individual embedded code signals. This entails the storage of a significant amount of additional code information above and beyond the original, and many applications may merit some form of economizing.

30 One such approach to economizing is to have a given set of individual embedded code signals be common to a batch of source materials. For example, one thousand images can all utilize the same basic set of individual embedded code signals. The storage requirements of these codes then become a small fraction of the overall storage requirements of the source material.

Furthermore, some applications can utilize a universal set of individual embedded
35 code signals, i.e., codes which remain the same for all instances of distributed material. This type of requirement would be seen by systems which wish to hide the N-bit identification word itself, yet have standardized equipment be able to read that word. This can be used in systems which make go/no go decisions at point-of-read locations. The potential drawback to this set-up is that

the universal codes are more prone to be sleuthed or stolen; therefore they will not be as secure as the apparatus and methodology of the previously disclosed arrangement. Perhaps this is just the difference between 'high security' and 'air-tight security,' a distinction carrying little weight with the bulk of potential applications.

5 Use in Printing, Paper, Documents, Plastic Coated Identification Cards, and Other Material Where Global Embedded Codes Can Be Imprinted

The term 'signal' is often used narrowly to refer to digital data information, audio signals, images, etc. A broader interpretation of 'signal,' and the one more generally intended, includes any form of modulation of any material whatsoever. Thus, the micro-topology of a piece of common paper becomes a 'signal' (e.g. its height as a function of x-y coordinates).
10 The reflective properties of a flat piece of plastic (as a function of space also) becomes a signal. The point is that photographic emulsions, audio signals, and digitized information are not the only types of signals capable of utilizing the principles of the present invention.

As a case in point, a machine very much resembling a braille printing machine
15 can be designed so as to imprint unique 'noise-like' indentations as outlined above. These indentations can be applied with a pressure which is much smaller than is typically applied in creating braille, to the point where the patterns are not noticed by a normal user of the paper. But by following the steps of the present disclosure and applying them via the mechanism of micro-indentations, a unique identification code can be placed onto any given sheet of paper, be it
20 intended for everyday stationary purposes, or be it for important documents, legal tender, or other secured material.

The reading of the identification material in such an embodiment generally proceeds by merely reading the document optically at a variety of angles. This would become an inexpensive method for deducing the micro-topology of the paper surface. Certainly other forms
25 of reading the topology of the paper are possible as well.

In the case of plastic encased material such as identification cards, e.g. driver's licenses, a similar braille-like impressions machine can be utilized to imprint unique identification codes. Subtle layers of photoreactive materials can also be embedded inside the plastic and 'exposed.'

30 It is clear that wherever a material exists which is capable of being modulated by 'noise-like' signals, that material is an appropriate carrier for unique identification codes and utilization of the principles of the invention. All that remains is the matter of economically applying the identification information and maintaining the signal level below an acceptability threshold which each and every application will define for itself.

35 Appendix A Description

Appendix A contains the source code of an implementation and verification of the foregoing embodiment for an 8-bit black and white imaging system.

REAL TIME ENCODER

While the first class of embodiments most commonly employs a standard microprocessor or computer to perform the encodation of an image or signal, it is possible to utilize a custom encodation device which may be faster than a typical Von Neuman-type processor. Such a system can be utilized with all manner of serial data streams.

Music and videotape recordings are examples of serial data streams -- data streams which are often pirated. It would assist enforcement efforts if authorized recordings were encoded with identification data so that pirated knock-offs could be traced to the original from which they were made.

Piracy is but one concern driving the need for the present invention. Another is authentication. Often it is important to confirm that a given set of data is really what it is purported to be (often several years after its generation).

To address these and other needs, the system 200 of Fig. 5 can be employed. System 200 can be thought of as an identification coding black box 202. The system 200 receives an input signal (sometimes termed the "master" or "unencoded" signal) and a code word, and produces (generally in real time) an identification-coded output signal. (Usually, the system provides key data for use in later decoding.)

The contents of the "black box" 202 can take various forms. An exemplary black box system is shown in Fig. 6 and includes a look-up table 204, a digital noise source 206, first and second scalers 208, 210, an adder/subtractor 212, a memory 214, and a register 216.

The input signal (which in the illustrated embodiment is an 8 - 20 bit data signal provided at a rate of one million samples per second, but which in other embodiments could be an analog signal if appropriate A/D and D/A conversion is provided) is applied from an input 218 to the address input 220 of the look-up table 204. For each input sample (i.e. look-up table address), the table provides a corresponding 8-bit digital output word. This output word is used as a scaling factor that is applied to one input of the first scaler 208.

The first scaler 208 has a second input, to which is applied an 8-bit digital noise signal from source 206. (In the illustrated embodiment, the noise source 206 comprises an analog noise source 222 and an analog-to-digital converter 224 although, again, other implementations can be used.) The noise source in the illustrated embodiment has a zero mean output value, with a full width half maximum (FWHM) of 50 - 100 digital numbers (e.g. from -75 to +75).

The first scaler 208 multiplies the two 8-bit words at its inputs (scale factor and noise) to produce -- for each sample of the system input signal -- a 16-bit output word. Since the noise signal has a zero mean value, the output of the first scaler likewise has a zero mean value.

The output of the first scaler 208 is applied to the input of the second scaler 210. The second scaler serves a global scaling function, establishing the absolute magnitude of the identification signal that will ultimately be embedded into the input data signal. The scaling factor is set through a scale control device 226 (which may take a number of forms, from a simple

rheostat to a graphically implemented control in a graphical user interface), permitting this factor to be changed in accordance with the requirements of different applications. The second scaler 210 provides on its output line 228 a scaled noise signal. Each sample of this scaled noise signal is successively stored in the memory 214.

5 (In the illustrated embodiment, the output from the first scaler 208 may range between -1500 and +1500 (decimal), while the output from the second scaler 210 is in the low single digits, (such as between -2 and +2).)

10 Register 216 stores a multi-bit identification code word. In the illustrated embodiment this code word consists of 8 bits, although larger code words (up to hundreds of bits) are commonly used. These bits are referenced, one at a time, to control how the input signal is modulated with the scaled noise signal.

15 In particular, a pointer 230 is cycled sequentially through the bit positions of the code word in register 216 to provide a control bit of "0" or "1" to a control input 232 of the adder/subtractor 212. If, for a particular input signal sample, the control bit is a "1", the scaled noise signal sample on line 232 is added to the input signal sample. If the control bit is a "0", the scaled noise signal sample is subtracted from the input signal sample. The output 234 from the adder/subtractor 212 provides the black box's output signal.

20 The addition or subtraction of the scaled noise signal in accordance with the bits of the code word effects a modulation of the input signal that is generally imperceptible. However, with knowledge of the contents of the memory 214, a user can later decode the encoding, determining the code number used in the original encoding process. (Actually, use of memory 214 is optional, as explained below.)

25 It will be recognized that the encoded signal can be distributed in well known ways, including converted to printed image form, stored on magnetic media (floppy diskette, analog or DAT tape, etc.), CD-ROM, etc. etc.

Decoding

30 A variety of techniques can be used to determine the identification code with which a suspect signal has been encoded. Two are discussed below. The first is less preferable than the latter for most applications, but is discussed herein so that the reader may have a fuller context within which to understand the invention.

More particularly, the first decoding method is a difference method, relying on subtraction of corresponding samples of the original signal from the suspect signal to obtain difference samples, which are then examined (typically individually) for deterministic coding indicia (i.e. the stored noise data). This approach may thus be termed a "sample-based, deterministic" decoding technique.

35 The second decoding method does not make use of the original signal. Nor does it examine particular samples looking for predetermined noise characteristics. Rather, the statistics of the suspect signal (or a portion thereof) are considered in the aggregate and analyzed to discern

the presence of identification coding that permeates the entire signal. The reference to permeation means the entire identification code can be discerned from a small fragment of the suspect signal. This latter approach may thus be termed a "holographic, statistical" decoding technique.

Both of these methods begin by registering the suspect signal to match the
5 original. This entails scaling (e.g. in amplitude, duration, color balance, etc.), and sampling (or resampling) to restore the original sample rate. As in the earlier described embodiment, there are a variety of well understood techniques by which the operations associated with this registration function can be performed.

As noted, the first decoding approach proceeds by subtracting the original signal
10 from the registered, suspect signal, leaving a difference signal. The polarity of successive difference signal samples can then be compared with the polarities of the corresponding stored noise signal samples to determine the identification code. That is, if the polarity of the first difference signal sample matches that of the first noise signal sample, then the first bit of the identification code is a "1." (In such case, the polarity of the 9th, 17th, 25th, etc. samples should also all be positive.) If the polarity of the first difference signal sample is opposite that of the
15 corresponding noise signal sample, then the first bit of the identification code is a "0."

By conducting the foregoing analysis with eight successive samples of the difference signal, the sequence of bits that comprise the original code word can be determined. If, as in the preferred embodiment, pointer 230 stepped through the code word one bit at a time,
20 beginning with the first bit, during encoding, then the first 8 samples of the difference signal can be analyzed to uniquely determine the value of the 8-bit code word.

In a noise-free world (speaking here of noise independent of that with which the identification coding is effected), the foregoing analysis would always yield the correct identification code. But a process that is only applicable in a noise-free world is of limited utility indeed.
25

(Further, accurate identification of signals in noise-free contexts can be handled in a variety of other, simpler ways: e.g. checksums; statistically improbable correspondence between suspect and original signals; etc.)

While noise-induced aberrations in decoding can be dealt with -- to some degree
30 -- by analyzing large portions of the signal, such aberrations still place a practical ceiling on the confidence of the process. Further, the villain that must be confronted is not always as benign as random noise. Rather, it increasingly takes the form of human-caused corruption, distortion, manipulation, etc. In such cases, the desired degree of identification confidence can only be achieved by other approaches.

35 The presently preferred approach (the "holographic, statistical" decoding technique) relies on recombining the suspect signal with certain noise data (typically the data stored in memory 214), and analyzing the entropy of the resulting signal. "Entropy" need not be

-23-

understood in its most strict mathematical definition, it being merely the most concise word to describe randomness (noise, smoothness, snowiness, etc.).

Most serial data signals are not random. That is, one sample usually correlates -- to some degree -- with the adjacent samples. Noise, in contrast, typically is random. If a random 5 signal (e.g. noise) is added to (or subtracted from) a non-random signal, the entropy of the resulting signal generally increases. That is, the resulting signal has more random variations than the original signal. This is the case with the encoded output signal produced by the present encoding process; it has more entropy than the original, unencoded signal.

If, in contrast, the addition of a random signal to (or subtraction from) a non- 10 random signal reduces entropy, then something unusual is happening. It is this anomaly that the preferred decoding process uses to detect embedded identification coding.

To fully understand this entropy-based decoding method, it is first helpful to highlight a characteristic of the original encoding process: the similar treatment of every eighth sample.

15 In the encoding process discussed above, the pointer 230 increments through the code word, one bit for each successive sample of the input signal. If the code word is eight bits in length, then the pointer returns to the same bit position in the code word every eighth signal sample. If this bit is a "1", noise is added to the input signal; if this bit is a "0", noise is subtracted from the input signal. Due to the cyclic progression of the pointer 230, every eighth 20 sample of an encoded signal thus shares a characteristic: they are all either augmented by the corresponding noise data (which may be negative), or they are all diminished, depending on whether the bit of the code word then being addressed by pointer 230 is a "1" or a "0".

25 To exploit this characteristic, the entropy-based decoding process treats every eighth sample of the suspect signal in like fashion. In particular, the process begins by adding to the 1st, 9th, 17th, 25th, etc. samples of the suspect signal the corresponding scaled noise signal values stored in the memory 214 (i.e. those stored in the 1st, 9th, 17th, 25th, etc., memory locations, respectively). The entropy of the resulting signal (i.e. the suspect signal with every 8th sample modified) is then computed.

30 (Computation of a signal's entropy or randomness is well understood by artisans in this field. One generally accepted technique is to take the derivative of the signal at each sample point, square these values, and then sum over the entire signal. However, a variety of other well known techniques can alternatively be used.)

The foregoing step is then repeated, this time subtracting the stored noise values from the 1st, 9th, 17th, 25 etc. suspect signal samples.

35 One of these two operations will undo the encoding process and reduce the resulting signal's entropy; the other will aggravate it. If adding the noise data in memory 214 to the suspect signal reduces its entropy, then this data must earlier have been subtracted from the original signal. This indicates that pointer 230 was pointing to a "0" bit when these samples were

encoded. (A "0" at the control input of adder/subtractor 212 caused it to subtract the scaled noise from the input signal.)

Conversely, if subtracting the noise data from every eighth sample of the suspect signal reduces its entropy, then the encoding process must have earlier added this noise. This
5 indicates that pointer 230 was pointing to a "1" bit when samples 1, 9, 17, 25, etc., were encoded.

By noting whether entropy decreases by (a) adding or (b) subtracting the stored noise data to/from the suspect signal, it can be determined that the first bit of the code word is (a) a "0", or (b) a "1".

10 The foregoing operations are then conducted for the group of spaced samples of the suspect signal beginning with the second sample (i.e. 2, 10, 18, 26 ...). The entropy of the resulting signals indicate whether the second bit of the code word is a "0" or a "1". Likewise with the following 6 groups of spaced samples in the suspect signal, until all 8 bits of the code word have been discerned.

15 It will be appreciated that the foregoing approach is not sensitive to corruption mechanisms that alter the values of individual samples; instead, the process considers the entropy of the signal as a whole, yielding a high degree of confidence in the results. Further, even small excerpts of the signal can be analyzed in this manner, permitting piracy of even small details of an original work to be detected. The results are thus statistically robust, both in the face of natural and human corruption of the suspect signal.

20 It will further be appreciated that the use of an N-bit code word in this real time embodiment provides benefits analogous to those discussed above in connection with the batch encoding system. (Indeed, the present embodiment may be conceptualized as making use of N different noise signals, just as in the batch encoding system. The first noise signal is a signal having the same extent as the input signal, and comprising the scaled noise signal at the 1st, 9th, 25 17th, 25th, etc., samples (assuming N=8), with zeroes at the intervening samples. The second noise signal is a similar one comprising the scaled noise signal at the 2d, 10th, 18th, 26th, etc., samples, with zeroes at the intervening samples. Etc. These signals are all combined to provide a composite noise signal.) One of the important advantages inherent in such a system is the high degree of statistical confidence (confidence which doubles with each successive bit of the
30 identification code) that a match is really a match. The system does not rely on subjective evaluation of a suspect signal for a single, deterministic embedded code signal.

Illustrative Variations

From the foregoing description, it will be recognized that numerous modifications can be made to the illustrated systems without changing the fundamental principles.
35 A few of these variations are described below.

The above-described decoding process tries both adding and subtracting stored noise data to/from the suspect signal in order to find which operation reduces entropy. In other embodiments, only one of these operations needs to be conducted. For example, in one alternative

decoding process the stored noise data corresponding to every eighth sample of the suspect signal is only added to said samples. If the entropy of the resulting signal is thereby increased, then the corresponding bit of the code word is a "1" (i.e. this noise was added earlier, during the encoding process, so adding it again only compounds the signal's randomness). If the entropy of the 5 resulting signal is thereby decreased, then the corresponding bit of the code word is a "0". A further test of entropy if the stored noise samples are subtracted is not required.

The statistical reliability of the identification process (coding and decoding) can be designed to exceed virtually any confidence threshold (e.g. 99.9%, 99.99%, 99.999%, etc. 10 confidence) by appropriate selection of the global scaling factors, etc. Additional confidence in any given application (unnecessary in most applications) can be achieved by rechecking the decoding process.

One way to recheck the decoding process is to remove the stored noise data from the suspect signal in accordance with the bits of the discerned code word, yielding a "restored" 15 signal (e.g. if the first bit of the code word is found to be "1," then the noise samples stored in the 1st, 9th, 17th, etc. locations of the memory 214 are subtracted from the corresponding samples of the suspect signal). The entropy of the restored signal is measured and used as a baseline in further measurements. Next, the process is repeated, this time removing the stored noise data from the suspect signal in accordance with a modified code word. The modified code word is the same as the discerned code word, except 1 bit is toggled (e.g. the first). The entropy of the 20 resulting signal is determined, and compared with the baseline. If the toggling of the bit in the discerned code word resulted in increased entropy, then the accuracy of that bit of the discerned code word is confirmed. The process repeats, each time with a different bit of the discerned code word toggled, until all bits of the code word have been so checked. Each change should result in an increase in entropy compared to the baseline value.

25 The data stored in memory 214 is subject to a variety of alternatives. In the foregoing discussion, memory 214 contains the scaled noise data. In other embodiments, the unscaled noise data can be stored instead.

In still other embodiments, it can be desirable to store at least part of the input 30 signal itself in memory 214. For example, the memory can allocate 8 signed bits to the noise sample, and 16 bits to store the most significant bits of an 18- or 20-bit audio signal sample. This has several benefits. One is that it simplifies registration of a "suspect" signal. Another is that, in the case of encoding an input signal which was already encoded, the data in memory 214 can be used to discern which of the encoding processes was performed first. That is, from the input signal data in memory 214 (albeit incomplete), it is generally possible to determine with which of 35 two code words it has been encoded.

Yet another alternative for memory 214 is that it can be omitted altogether.

One way this can be achieved is to use a deterministic noise source in the encoding process, such as an algorithmic noise generator seeded with a known key number. The

same deterministic noise source, seeded with the same key number, can be used in the decoding process. In such an arrangement, only the key number needs be stored for later use in decoding, instead of the large data set usually stored in memory 214.

Alternatively, if the noise signal added during encoding does not have a zero
5 mean value, and the length N of the code word is known to the decoder, then a universal decoding process can be implemented. This process uses the same entropy test as the foregoing procedures, but cycles through possible code words, adding/subtracting a small dummy noise value (e.g. less than the expected mean noise value) to every Nth sample of the suspect signal, in accordance with the bits of the code word being tested, until a reduction in entropy is noted. Such an approach is
10 not favored for most applications, however, because it offers less security than the other embodiments (e.g. it is subject to cracking by brute force).

Many applications are well served by the embodiment illustrated in Fig. 7, in which different code words are used to produce several differently encoded versions of an input signal, each making use of the same noise data. More particularly, the embodiment 240 of Fig. 7
15 includes a noise store 242 into which noise from source 206 is written during the identification-coding of the input signal with a first code word. (The noise source of Fig. 7 is shown outside of the real time encoder 202 for convenience of illustration.) Thereafter, additional identification-coded versions of the input signal can be produced by reading the stored noise data from the store and using it in conjunction with second through Nth code words to encode the signal. (While
20 binary-sequential code words are illustrated in Fig. 7, in other embodiments arbitrary sequences of code words can be employed.) With such an arrangement, a great number of differently-encoded signals can be produced, without requiring a proportionally-sized long term noise memory. Instead, a fixed amount of noise data is stored, whether encoding an original once or a thousand times.

25 (If desired, several differently-coded output signals can be produced at the same time, rather than seriatim. One such implementation includes a plurality of adder/subtractor circuits 212, each driven with the same input signal and with the same scaled noise signal, but with different code words. Each, then, produces a differently encoded output signal.)

In applications having a great number of differently-encoded versions of the
30 same original, it will be recognized that the decoding process need not always discern every bit of the code word. Sometimes, for example, the application may require identifying only a group of codes to which the suspect signal belongs. (E.g., high order bits of the code word might indicate an organization to which several differently coded versions of the same source material were provided, with low-order bits identifying specific copies. To identify the organization with which
35 a suspect signal is associated, it may not be necessary to examine the low order bits, since the organization can be identified by the high order bits alone.) If the identification requirements can be met by discerning a subset of the code word bits in the suspect signal, the decoding process can be shortened.

Some applications may be best served by restarting the encoding process -- sometimes with a different code word -- several times within an integral work. Consider, as an example, videotaped productions (e.g. television programming). Each frame of a videotaped production can be identification-coded with a unique code number, processed in real-time with an arrangement 248 like that shown in Fig. 8. Each time a vertical retrace is detected by sync detector 250, the noise source 206 resets (e.g. to repeat the sequence just produced) and an identification code increments to the next value. Each frame of the videotape is thereby uniquely identification-coded. Typically, the encoded signal is stored on a videotape for long term storage (although other storage media, including laser disks, can be used).

Returning to the encoding apparatus, the look-up table 204 in the illustrated embodiment exploits the fact that high amplitude samples of the input data signal can tolerate (without objectionable degradation of the output signal) a higher level of encoded identification coding than can low amplitude input samples. Thus, for example, input data samples having decimal values of 0, 1 or 2 may correspond (in the look-up table 204) to scale factors of unity (or even zero), whereas input data samples having values in excess of 200 may correspond to scale factors of 15. Generally speaking, the scale factors and the input sample values correspond by a square root relation. That is, a four-fold increase in a value of the sampled input signal corresponds to approximately a two-fold increase in a value of the scaling factor associated therewith.

(The parenthetical reference to zero as a scaling factor alludes to cases, e.g., in which the source signal is temporally or spatially devoid of information content. In an image, for example, a region characterized by several contiguous sample values of zero may correspond to a jet black region of the frame. A scaling value of zero may be appropriate here since there is essentially no image data to be pirated.)

Continuing with the encoding process, those skilled in the art will recognize the potential for "rail errors" in the illustrated embodiment. For example, if the input signal consists of 8-bit samples, and the samples span the entire range from 0 to 255 (decimal), then the addition or subtraction of scaled noise to/from the input signal may produce output signals that cannot be represented by 8 bits (e.g. -2, or 257). A number of well-understood techniques exist to rectify this situation, some of them proactive and some of them reactive. (Among these known techniques are: specifying that the input signal shall not have samples in the range of 0-4 or 251-255, thereby safely permitting modulation by the noise signal; or including provision for detecting and adaptively modifying input signal samples that would otherwise cause rail errors.)

While the illustrated embodiment describes stepping through the code word sequentially, one bit at a time, to control modulation of successive bits of the input signal, it will be appreciated that the bits of the code word can be used other than sequentially for this purpose. Indeed, bits of the code word can be selected in accordance with any predetermined algorithm.

The dynamic scaling of the noise signal based on the instantaneous value of the input signal is an optimization that can be omitted in many embodiments. That is, the look-up table 204 and the first scaler 208 can be omitted entirely, and the signal from the digital noise source 206 applied directly (or through the second, global scaler 210) to the adder/subtractor 212.

5 It will be further recognized that the use of a zero-mean noise source simplifies the illustrated embodiment, but is not necessary to the invention. A noise signal with another mean value can readily be used, and D.C. compensation (if needed) can be effected elsewhere in the system.

10 The use of a noise source 206 is also optional. A variety of other signal sources can be used, depending on application-dependent constraints (e.g. the threshold at which the encoded identification signal becomes perceptible). In many instances, the level of the embedded identification signal is low enough that the identification signal needn't have a random aspect; it is imperceptible regardless of its nature. A pseudo random source 206, however, is usually desired because it provides the greatest identification code signal S/N ratio (a somewhat awkward term in 15 this instance) for a level of imperceptibility of the embedded identification signal.

15 It will be recognized that identification coding need not occur after a signal has been reduced to stored form as data (i.e. "fixed in tangible form," in the words of the U.S. Copyright Act). Consider, for example, the case of popular musicians whose performances are often recorded illicitly. By identification coding the audio before it drives concert hall speakers, 20 unauthorized recordings of the concert can be traced to a particular place and time. Likewise, live audio sources such as 911 emergency calls can be encoded prior to recording so as to facilitate their later authentication.

25 While the black box embodiment has been described as a stand alone unit, it will be recognized that it can be integrated into a number of different tools/instruments as a component. One is a scanner, which can embed identification codes in the scanned output data. (The codes can simply serve to memorialize that the data was generated by a particular scanner). Another is in creativity software, such as popular drawing/graphics/animation/paint programs offered by Adobe, Macromedia, Corel, and the like.

30 Finally, while the real-time encoder 202 has been illustrated with reference to a particular hardware implementation, it will be recognized that a variety of other implementations can alternatively be employed. Some utilize other hardware configurations. Others make use of software routines for some or all of the illustrated functional blocks. (The software routines can be executed on any number of different general purpose programmable computers, such as 80x86 PC-compatible computers, RISC-based workstations, etc.)

35 TYPES OF NOISE, QUASI-NOISE, AND OPTIMIZED-NOISE

Heretofore this disclosure postulated Gaussian noise, "white noise," and noise generated directly from application instrumentation as a few of the many examples of the kind of carrier signal appropriate to carry a single bit of information throughout an image or signal. It is

-29-

possible to be even more proactive in "designing" characteristics of noise in order to achieve certain goals. The "design" of using Gaussian or instrumental noise was aimed somewhat toward "absolute" security. This section of the disclosure takes a look at other considerations for the design of the noise signals which may be considered the ultimate carriers of the identification information.

For some applications it might be advantageous to design the noise carrier signal (e.g. the Nth embedded code signal in the first embodiment; the scaled noise data in the second embodiment), so as to provide more absolute signal strength to the identification signal relative to the perceptibility of that signal. One example is the following. It is recognized that a true Gaussian noise signal has the value '0' occur most frequently, followed by 1 and -1 at equal probabilities to each other but lower than '0', 2 and -2 next, and so on. Clearly, the value zero carries no information as it is used in the service of this invention. Thus, one simple adjustment, or design, would be that any time a zero occurs in the generation of the embedded code signal, a new process takes over, whereby the value is converted "randomly" to either a 1 or a -1. In logical terms, a decision would be made: if '0', then random(1,-1). The histogram of such a process would appear as a Gaussian/Poissonian type distribution, except that the 0 bin would be empty and the 1 and -1 bin would be increased by half the usual histogram value of the 0 bin.

In this case, identification signal energy would always be applied at all parts of the signal. A few of the trade-offs include: there is a (probably negligible) lowering of security of the codes in that a "deterministic component" is a part of generating the noise signal. The reason this might be completely negligible is that we still wind up with a coin flip type situation on randomly choosing the 1 or the -1. Another trade-off is that this type of designed noise will have a higher threshold of perceptibility, and will only be applicable to applications where the least significant bit of a data stream or image is already negligible relative to the commercial value of the material, i.e. if the least significant bit were stripped from the signal (for all signal samples), no one would know the difference and the value of the material would not suffer. This blocking of the zero value in the example above is but one of many ways to "optimize" the noise properties of the signal carrier, as anyone in the art can realize. We refer to this also as "quasi-noise" in the sense that natural noise can be transformed in a pre-determined way into signals which for all intents and purposes will read as noise. Also, cryptographic methods and algorithms can easily, and often by definition, create signals which are perceived as completely random. Thus the word "noise" can have different connotations, primarily between that as defined subjectively by an observer or listener, and that defined mathematically. The difference of the latter is that mathematical noise has different properties of security and the simplicity with which it can either be "sleuthed" or the simplicity with which instruments can "automatically recognize" the existence of this noise.

"Universal" Embedded Codes

- The bulk of this disclosure teaches that for absolute security, the noise-like embedded code signals which carry the bits of information of the identification signal should be unique to each and every encoded signal, or, slightly less restrictive, that embedded code signals should be generated sparingly, such as using the same embedded codes for a batch of 1000 pieces of film, for example. Be this as it may, there is a whole other approach to this issue wherein the use of what we will call "universal" embedded code signals can open up large new applications for this technology. The economics of these uses would be such that the de facto lowered security of these universal codes (e.g. they would be analyzable by time honored cryptographic decoding methods, and thus potentially thwarted or reversed) would be economically negligible relative to the economic gains that the intended uses would provide. Piracy and illegitimate uses would become merely a predictable "cost" and a source of uncollected revenue only; a simple line item in an economic analysis of the whole. A good analogy of this is in the cable industry and the scrambling of video signals. Everybody seems to know that crafty, skilled technical individuals, who may be generally law abiding citizens, can climb a ladder and flip a few wires in their cable junction box in order to get all the pay channels for free. The cable industry knows this and takes active measures to stop it and prosecute those caught, but the "lost revenue" derived from this practice remains prevalent but almost negligible as a percentage of profits gained from the scrambling system as a whole. The scrambling system as a whole is an economic success despite its lack of "absolute security."
- The same holds true for applications of this technology wherein, for the price of lowering security by some amount, large economic opportunity presents itself. This section first describes what is meant by universal codes, then moves on to some of the interesting uses to which these codes can be applied.
- Universal embedded codes generally refer to the idea that knowledge of the exact codes can be distributed. The embedded codes won't be put into a dark safe never to be touched until litigation arises (as alluded to in other parts of this disclosure), but instead will be distributed to various locations where on-the-spot analysis can take place. Generally this distribution will still take place within a security controlled environment, meaning that steps will be taken to limit the knowledge of the codes to those with a need to know. Instrumentation which attempts to automatically detect copyrighted material is a non-human example of "something" with a need to know the codes.
- There are many ways to implement the idea of universal codes, each with their own merits regarding any given application. For the purposes of teaching this art, we separate these approaches into three broad categories: universal codes based on libraries, universal codes based on deterministic formula, and universal codes based on pre-defined industry standard patterns. A rough rule of thumb is that the first is more secure than the latter two, but that the latter two are possibly more economical to implement than the first.

Universal Codes: 1) Libraries of Universal Codes

The use of libraries of universal codes simply means that the techniques of this invention are employed as described, except for the fact that only a limited set of the individual embedded code signals are generated and that any given encoded material will make use of some 5 sub-set of this limited "universal set." An example is in order here. A photographic print paper manufacturer may wish to pre-expose every piece of 8 by 10 inch print paper which they sell with a unique identification code. They also wish to sell identification code recognition software to their large customers, service bureaus, stock agencies, and individual photographers, so that all these people can not only verify that their own material is correctly marked, but so that they can 10 also determine if third party material which they are about to acquire has been identified by this technology as being copyrighted. This latter information will help them verify copyright holders and avoid litigation, among many other benefits. In order to "economically" institute this plan, they realize that generating unique individual embedded codes for each and every piece of print paper would generate Terabytes of independent information, which would need storing and to 15 which recognition software would need access. Instead, they decide to embed their print paper with 16 bit identification codes derived from a set of only 50 independent "universal" embedded code signals. The details of how this is done are in the next paragraph, but the point is that now their recognition software only needs to contain a limited set of embedded codes in their library of codes, typically on the order of 1 Megabyte to 10 Megabytes of information for 50x16 individual 20 embedded codes splayed out onto an 8x10 photographic print (allowing for digital compression). The reason for picking 50 instead of just 16 is one of a little more added security, where if it were the same 16 embedded codes for all photographic sheets, not only would the serial number capability be limited to 2 to the 16th power, but lesser and lesser sophisticated pirates could crack the codes and remove them using software tools.

25 There are many different ways to implement this scheme, where the following is but one exemplary method. It is determined by the wisdom of company management that a 300 pixels per inch criteria for the embedded code signals is sufficient resolution for most applications. This means that a composite embedded code image will contain 3000 pixels by 2400 pixels to be exposed at a very low level onto each 8x10 sheet. This gives 7.2 million pixels. Using our 30 staggered coding system such as described in the black box implementation of Figs. 5 and 6, each individual embedded code signal will contain only 7.2 million divided by 16, or approximately 450K true information carrying pixels, i.e. every 16th pixel along a given raster line. These values will typically be in the range of 2 to -2 in digital numbers, or adequately described by a signed 3 bit number. The raw information content of an embedded code is then approximately 35 3/8th's bytes times 450K or about 170 Kilobytes. Digital compression can reduce this further. All of these decisions are subject to standard engineering optimization principles as defined by any given application at hand, as is well known in the art. Thus we find that 50 of these independent embedded codes will amount to a few Megabytes. This is quite reasonable level to distribute as a

-32-

"library" of universal codes within the recognition software. Advanced standard encryption devices could be employed to mask the exact nature of these codes if one were concerned that would-be pirates would buy the recognition software merely to reverse engineer the universal embedded codes. The recognition software could simply unencrypt the codes prior to applying the 5 recognition techniques taught in this disclosure.

The recognition software itself would certainly have a variety of features, but the core task it would perform is determining if there is some universal copyright code within a given image. The key questions become WHICH 16 of the total 50 universal codes it might contain, if any, and if there are 16 found, what are their bit values. The key variables in determining the 10 answers to these questions are: registration, rotation, magnification (scale), and extent. In the most general case with no helpful hints whatsoever, all variables must be independently varied across all mutual combinations, and each of the 50 universal codes must then be checked by adding and subtracting to see if an entropy decrease occurs. Strictly speaking, this is an enormous job, but many helpful hints will be found which make the job much simpler, such as having an original 15 image to compare to the suspected copy, or knowing the general orientation and extent of the image relative to an 8x10 print paper, which then through simple registration techniques can determine all of the variables to some acceptable degree. Then it merely requires cycling through the 50 universal codes to find any decrease in entropy. If one does, then 15 others should as well. A protocol needs to be set up whereby a given order of the 50 translates into a sequence of most 20 significant bit through least significant bit of the ID code word. Thus if we find that universal code number "4" is present, and we find its bit value to be "0", and that universal codes "1" through "3" are definitely not present, then our most significant bit of our N-bit ID code number is a "0". Likewise, we find that the next lowest universal code present is number "7" and it turns 25 out to be a "1", then our next most significant bit is a "1". Done properly, this system can cleanly trace back to the copyright owner so long as they registered their photographic paper stock serial number with some registry or with the manufacturer of the paper itself. That is, we look up in the registry that a paper using universal embedded codes 4,7,11,12,15,19,21,26,27,28,34,35,37,38,40, and 48, and having the embedded code 0110 0101 0111 0100 belongs to Leonardo de Boticelli, an unknown wildlife photographer and glacier cinematographer whose address is in Northern Canada. 30 We know this because he dutifully registered his film and paper stock, a few minutes of work when he bought the stock, which he plopped into the "no postage necessary" envelope that the manufacturing company kindly provided to make the process ridiculously simple. Somebody owes Leonardo a royalty check it would appear, and certainly the registry has automated this royalty payment process as part of its services.

35 One final point is that truly sophisticated pirates and others with illicit intentions can indeed employ a variety of cryptographic and not so cryptographic methods to crack these universal codes, sell them, and make software and hardware tools which can assist in the removing or distorting of codes. We shall not teach these methods as part of this disclosure, however. In

-33-

any event, this is one of the prices which must be paid for the ease of universal codes and the applications they open up.

Universal Codes: 2) Universal Codes Based on Deterministic Formulas

The libraries of universal codes require the storage and transmittal of Megabytes of independent, generally random data as the keys with which to unlock the existence and identity of signals and imagery that have been marked with universal codes. Alternatively, various deterministic formulas can be used which "generate" what appear to be random data/image frames, thereby obviating the need to store all of these codes in memory and interrogate each and of the "50" universal codes. Deterministic formulas can also assist in speeding up the process of determining the ID code once one is known to exist in a given signal or image. On the other hand, deterministic formulas lend themselves to sleuthing by less sophisticated pirates. And once sleuthed, they lend themselves to easier communication, such as posting on the Internet to a hundred newsgroups. There may well be many applications which do not care about sleuthing and publishing, and deterministic formulas for generating the individual universal embedded codes might be just the ticket.

Universal Codes: 3) "Simple" Universal Codes

This category is a bit of a hybrid of the first two, and is most directed at truly large scale implementations of the principles of this technology. The applications employing this class are of the type where staunch security is much less important than low cost, large scale implementation and the vastly larger economic benefits that this enables. One exemplary application is placement of identification recognition units directly within modestly priced home audio and video instrumentation (such as a TV). Such recognition units would typically monitor audio and/or video looking for these copyright identification codes, and thence triggering simple decisions based on the findings, such as disabling or enabling recording capabilities, or incrementing program specific billing meters which are transmitted back to a central audio/video service provider and placed onto monthly invoices. Likewise, it can be foreseen that "black boxes" in bars and other public places can monitor (listen with a microphone) for copyrighted materials and generate detailed reports, for use by ASCAP, BMI, and the like.

A core principle of simple universal codes is that some basic industry standard "noiselike" and seamlessly repetitive patterns are injected into signals, images, and image sequences so that inexpensive recognition units can either A) determine the mere existence of a copyright "flag", and B) additionally to A, determine precise identification information which can facilitate more complex decision making and actions.

In order to implement this particular embodiment of the present invention, the basic principles of generating the individual embedded noise signals need to be simplified in order to accommodate inexpensive recognition signal processing circuitry, while maintaining the properties of effective randomness and holographic permeation. With large scale industry adoption of these simple codes, the codes themselves would border on public domain information (much as

cable scrambling boxes are almost de facto public domain), leaving the door open for determined pirates to develop black market countermeasures, but this situation would be quite analogous to the scrambling of cable video and the objective economic analysis of such illegal activity.

One prior art known to the applicant in this general area of pro-active copyright detection is the Serial Copy Management System adopted by many firms in the audio industry. To the best of applicant's knowledge, this system employs a non-audio "flag" signal which is not part of the audio data stream, but which is nevertheless grafted onto the audio stream and can indicate whether the associated audio data should or should not be duplicated. One problem with this system is that it is restricted to media and instrumentation which can support this extra "flag" signal. Another deficiency is that the flagging system carries no identity information which would be useful in making more complex decisions. Yet another difficulty is that high quality audio sampling of an analog signal can come arbitrarily close to making a perfect digital copy of some digital master and there seems to be no provision for inhibiting this possibility.

The principles of this invention can be brought to bear on these and other problems, in audio applications, video, and all of the other applications previously discussed. An exemplary application of simple universal codes is the following. A single industry standard "1.000000 second of noise" would be defined as the most basic indicator of the presence or absence of the copyright marking of any given audio signal. Fig. 9 has an example of what the waveform of an industry standard noise second might look like, both in the time domain 400 and the frequency domain 402. It is by definition a continuous function and would adapt to any combination of sampling rates and bit quantizations. It has a normalized amplitude and can be scaled arbitrarily to any digital signal amplitude. The signal level and the first M'th derivatives of the signal are continuous at the two boundaries 404 (Fig. 9C), such that when it is repeated, the "break" in the signal would not be visible (as a waveform) or audible when played through a high end audio system. The choice of 1 second is arbitrary in this example, where the precise length of the interval will be derived from considerations such as audibility, quasi-white noise status, seamless repeatability, simplicity of recognition processing, and speed with which a copyright marking determination can be made. The injection of this repeated noise signal onto a signal or image (again, at levels below human perception) would indicate the presence of copyright material. This is essentially a one bit identification code, and the embedding of further identification information will be discussed later on in this section. The use of this identification technique can extend far beyond the low cost home implementations discussed here, where studios could use the technique, and monitoring stations could be set up which literally monitor hundreds of channels of information simultaneously, searching for marked data streams, and furthermore searching for the associated identity codes which could be tied in with billing networks and royalty tracking systems.

This basic, standardized noise signature is seamlessly repeated over and over again and added to audio signals which are to be marked with the base copyright identification.

Part of the reason for the word "simple" is seen here: clearly pirates will know about this industry standard signal, but their illicit uses derived from this knowledge, such as erasure or corruption, will be economically minuscule relative to the economic value of the overall technique to the mass market. For most high end audio this signal will be some 80 to 100 dB down from full scale, or even much further; each situation can choose its own levels though certainly there will be recommendations. The amplitude of the signal can be modulated according to the audio signal levels to which the noise signature is being applied, i.e. the amplitude can increase significantly when a drum beats, but not so dramatically as to become audible or objectionable. These measures merely assist the recognition circuitry to be described.

Recognition of the presence of this noise signature by low cost instrumentation can be effected in a variety of ways. One rests on basic modifications to the simple principles of audio signal power metering. Software recognition programs can also be written, and more sophisticated mathematical detection algorithms can be applied to audio in order to make higher confidence detection identifications. In such embodiments, detection of the copyright noise signature involves comparing the time averaged power level of an audio signal with the time averaged power level of that same audio signal which has had the noise signature subtracted from it. If the audio signal with the noise signature subtracted has a lower power level than the unchanged audio signal, then the copyright signature is present and some status flag to that effect needs to be set. The main engineering subtleties involved in making this comparison include: dealing with audio speed playback discrepancies (e.g. an instrument might be 0.5% "slow" relative to exactly one second intervals); and, dealing with the unknown phase of the one second noise signature within any given audio (basically, its "phase" can be anywhere from 0 to 1 seconds). Another subtlety, not so central as the above two but which nonetheless should be addressed, is that the recognition circuits should not subtract a higher amplitude of the noise signature than was originally embedded onto the audio signal. Fortunately this can be accomplished by merely subtracting only a small amplitude of the noise signal, and if the power level goes down, this is an indication of "heading toward a trough" in the power levels. Yet another related subtlety is that the power level changes will be very small relative to the overall power levels, and calculations generally will need to be done with appropriate bit precision, e.g. 32 bit value operations and accumulations on 16-20 bit audio in the calculations of time averaged power levels.

Clearly, designing and packaging this power level comparison processing circuitry for low cost applications is an engineering optimization task. One trade-off will be the accuracy of making an identification relative to the "short-cuts" which can be made to the circuitry in order to lower its cost and complexity. A preferred embodiment for the placement of this recognition circuitry inside of instrumentation is through a single programmable integrated circuit which is custom made for the task. Fig. 10 shows one such integrated circuit 506. Here the audio signal comes in, 500, either as a digital signal or as an analog signal to be digitized inside the IC 500, and the output is a flag 502 which is set to one level if the copyright noise signature is

found, and to another level if it is not found. Also depicted is the fact that the standardized noise signature waveform is stored in Read Only Memory, 504, inside the IC 506. There will be a slight time delay between the application of an audio signal to the IC 506 and the output of a valid flag 502, due to the need to monitor some finite portion of the audio before a recognition can place. In this case, there may need to be a "flag valid" output 508 where the IC informs the external world if it has had enough time to make a proper determination of the presence or absence of the copyright noise signature.

There are a wide variety of specific designs and philosophies of designs applied to accomplishing the basic function of the IC 506 of Fig. 10. Audio engineers and digital signal processing engineers are able to generate several fundamentally different designs. One such design is depicted in Fig. 11 by a process 599, which itself is subject to further engineering optimization as will be discussed. Fig. 11 depicts a flow chart for any of: an analog signal processing network, a digital signal processing network, or programming steps in a software program. We find an input signal 600 which along one path is applied to a time averaged power meter 602, and the resulting power output itself treated as a signal P_{sig} . To the upper right we find the standard noise signature 504 which will be read out at 125% of normal speed, 604, thus changing its pitch, giving the "pitch changed noise signal" 606. Then the input signal has this pitch changed noise signal subtracted in step 608, and this new signal is applied to the same form of time averaged power meter as in 602, here labelled 610. The output of this operation is also a time based signal here labelled as P_{ps} , 610. Step 612 then subtracts the power signal 602 from the power signal 610, giving an output difference signal P_{out} , 613. If the universal standard noise signature does indeed exist on the input audio signal 600, then case 2, 616, will be created wherein a beat signal 618 of approximately 4 second period will show up on the output signal 613, and it remains to detect this beat signal with a step such as in Fig. 12, 622. Case 1, 614, is a steady noisy signal which exhibits no periodic beating. 125% at step 604 is chosen arbitrarily here, where engineering considerations would determine an optimal value, leading to different beat signal frequencies 618. Whereas waiting 4 seconds in this example would be quite a while, especially if you would want to detect at least two or three beats, Fig. 12 outlines how the basic design of Fig. 11 could be repeated and operated upon various delayed versions of the input signal, delayed by something like 1/20th of a second, with 20 parallel circuits working in concert each on a segment of the audio delayed by 0.05 seconds from their neighbors. In this way, a beat signal will show up approximately every 1/5th of a second and will look like a travelling wave down the columns of beat detection circuits. The existence or absence of this travelling beat wave triggers the detection flag 502. Meanwhile, there would be an audio signal monitor 624 which would ensure that, for example, at least two seconds of audio has been heard before setting the flag valid signal 508.

Though the audio example was described above, it should be clear to anyone in the art that the same type of definition of some repetitive universal noise signal or image could be applied to the many other signals, images, pictures, and physical media already discussed.

The above case deals only with a single bit plane of information, i.e., the noise signature signal is either there (1) or it isn't (0). For many applications, it would be nice to detect serial number information as well, which could then be used for more complex decisions, or for logging information on billing statements or whatnot. The same principles as the above would apply, but now there would be N independent noise signatures as depicted in Fig. 9 instead one single such signature. Typically, one such signature would be the master upon which the mere existence of a copyright marking is detected, and this would have generally higher power than the others, and then the other lower power "identification" noise signatures would be embedded into audio. Recognition circuits, once having found the existence of the primary noise signature, would then step through the other N noise signatures applying the same steps as described above. Where a beat signal is detected, this indicates the bit value of '1', and where no beat signal is detected, this indicates a bit value of '0'. It might be typical that N will equal 32, that way 2^{32} number of identification codes are available to any given industry employing this invention.

Use of this Technology When the Length of the Identification Code is 1

The principles of this invention can obviously be applied in the case where only a single presence or absence of an identification signal -- a fingerprint if you will -- is used to provide confidence that some signal or image is copyrighted. The example above of the industry standard noise signature is one case in point. We no longer have the added confidence of the coin flip analogy, we no longer have tracking code capabilities or basic serial number capabilities, but many applications may not require these attributes and the added simplicity of a single fingerprint might outweigh these other attributes in any event.

The "Wallpaper" Analogy

The term "holographic" has been used in this disclosure to describe how an identification code number is distributed in a largely integral form throughout an encoded signal or image. This also refers to the idea that any given fragment of the signal or image contains the entire unique identification code number. As with physical implementations of holography, there are limitations on how small a fragment can become before one begins to lose this property, where the resolution limits of the holographic media are the main factor in this regard for holography itself. In the case of an uncorrupted distribution signal which has used the encoding device of figure 5, and which furthermore has used our "designed noise" of above wherein the zero's were randomly changed to a 1 or -1, then the extent of the fragment required is merely N contiguous samples in a signal or image raster line, where N is as defined previously being the length of our identification code number. This is an informational extreme; practical situations where noise and corruption are operative will require generally one, two or higher orders of magnitude more samples than this simple number N. Those skilled in the art will recognize that there are many

variables involved in pinning down precise statistics on the size of the smallest fragment with which an identification can be made.

For tutorial purposes, the applicant also uses the analogy that the unique identification code number is "wallpapered" across and image (or signal). That is, it is repeated over and over again all throughout an image. This repetition of the ID code number can be regular, as in the use of the encoder of figure 5, or random itself, where the bits in the ID code 216 of figure 6 are not stepped through in a normal repetitive fashion but rather are randomly selected on each sample, and the random selection stored along with the value of the output 228 itself. in any event, the information carrier of the ID code, the individual embedded code signal, does change across the image or signal. Thus as the wallpaper analogy summarizes: the ID code repeats itself over and over, but the patterns that each repetition imprints change randomly accordingly to a generally unsleuthable key.

Lossy Data Compression

As earlier mentioned, the identification coding of the preferred embodiment withstands lossy data compression, and subsequent decompression. Such compression is finding increasing use, particularly in contexts such as the mass distribution of digitized entertainment programming (movies, etc.).

While data encoded according to the preferred embodiment of the present invention can withstand all types of lossy compression known to applicant, those expected to be most commercially important are the CCITT G3, CCITT G4, JPEG, MPEG and JBIG compression/decompression standards. The CCITT standards are widely used in black-and-white document compression (e.g. facsimile and document-storage). JPEG is most widely used with still images. MPEG is most widely used with moving images. JBIG is a likely successor to the CCITT standards for use with black-and-white imagery. Such techniques are well known to those in the lossy data compression field; a good overview can be found in Pennebaker et al, *JPEG, Still Image Data Compression Standard*, Van Nostrand Reinhold, N.Y., 1993.

Towards Steganography Proper and the Use of this Technology in Passing More Complex Messages or Information

This disclosure concentrates on what above was called wallpapering a single identification code across an entire signal. This appears to be a desirable feature for many applications. However, there are other applications where it might be desirable to pass messages or to embed very long strings of pertinent identification information in signals and images. One of many such possible applications would be where a given signal or image is meant to be manipulated by several different groups, and that certain regions of an image are reserved for each group's identification and insertion of pertinent manipulation information.

In these cases, the code word 216 in figure 6 can actually change in some pre-defined manner as a function of signal or image position. For example, in an image, the code could change for each and every raster line of the digital image. It might be a 16 bit code word,

-39-

216, but each scan line would have a new code word, and thus a 480 scan line image could pass a
980 (480 x 2 bytes) byte message. A receiver of the message would need to have access to either
the noise signal stored in memory 214, or would have to know the universal code structure of the
noise codes if that method of coding was being used. To the best of applicant's knowledge, this is
5 a novel approach to the mature field of steganography.

In all three of the foregoing applications of universal codes, it will often be
desirable to append a short (perhaps 8- or 16-bit) private code, which users would keep in their
own secured places, in addition to the universal code. This affords the user a further modicum of
security against potential erasure of the universal codes by sophisticated pirates.

10 Conclusion

In view of the great number of different embodiments to which the principles of
my invention can be put, it should be recognized that the detailed embodiments are illustrative
only and should not be taken as limiting the scope of my invention. Rather, I claim as my
invention all such embodiments as may come within the scope and spirit of the following claims,
15 and equivalents thereto.

-40-

APPENDIX A

```

#include "main.h"

#define XDIM 512L
#define XDIMR 512
#define YDIM 480L
#define BITS 8
#define RMS_VAL 5.0
#define NUM_NOISY 16
#define NUM_DEMOS 3
#define GRAD_THRESHOLD 10

struct char_buf {
    char filename[80];
    FILE *fp;
    fpos_t fpos;
    char buf[XDIMR];
};

struct uchar_buf {
    char filename[80];
    FILE *fp;
    fpos_t fpos;
    unsigned char buf[XDIMR];
};

struct int_buf {
    char filename[80];
    FILE *fp;
    fpos_t fpos;
    int buf[XDIMR];
};

struct cortex_s {
    char filename[80];
    FILE *fp;
    fpos_t fpos;
    unsigned char buf[XDIMR];
};

struct uchar_buf test_image;
struct char_buf snow_composite;
struct uchar_buf distributed_image;
struct uchar_buf temp_image;
struct int_buf temp_wordbuffer;
struct int_buf temp_wordbuffer2;
struct uchar_buf snow_images;
struct cortex_s cortex;

int demo=0; /* which demo is being performed, see notes */
int our_code; /* id value embedded onto image */
int found_code=0; /* holder for found code*/

int waitvbb(void){
    while( _inp(PORT_BASE)&8) ;
    while( !( _inp(PORT_BASE)&8) );
    return(1);
}

int grabb(void){
    waitvbb();
    _outp(PORT_BASE+1,0);
    _outp(PORT_BASE,8);
    waitvbb();
    waitvbb();
    _outp(PORT_BASE,0x10);
    return(1);
}

```

```

}

int livee(void) {
    _outp(PORT_BASE, 0x00);
    return(1);
}

int live_video(void) {
    livee();
    return(1);
}

int freeze_frame(void) {
    grabb();
    return(1);
}

int grab_frame(struct uchar_buf *image) {
    long i;

    grabb();
    fsetpos(image->fp, &image->fpos );
    fsetpos(cortex.fp, &cortex.fpos );
    for(i=0;i<YDIM;i++){
        fread(cortex.buf,sizeof(unsigned char),XDIMR,cortex.fp);
        fwrite(cortex.buf,sizeof(unsigned char),XDIMR,image->fp);
    }
    livee();
    return(1);
}

int wait_vertical_blanks(int number){
    long i;
    for(i=0;i<number;i++) waitvbb();
    return(1);
}

int clear_char_image(struct char_buf *charbuffer) {
    long i,j;
    char *pchar;
    fpos_t tmp_fpos;

    fsetpos(charbuffer->fp, &charbuffer->fpos );
    for(i=0;i<YDIM;i++){
        fgetpos(charbuffer->fp, &tmp_fpos );
        pchar = charbuffer->buf;
        fread(charbuffer->buf,sizeof(char),XDIMR,charbuffer->fp);
        for(j=0;j<XDIM;j++) *(pchar++) = 0;
        fsetpos(charbuffer->fp, &tmp_fpos );
        fwrite(charbuffer->buf,sizeof(char),XDIMR,charbuffer->fp);
    }
    return(1);
}

int display_uchar(struct uchar_buf *image,int stretch) {
    unsigned char *pimage;
    unsigned char highest = 0;
    unsigned char lowest = 255;
    long i,j;
    double dtemp,scale,dlowest;
    fpos_t tmp_fpos;

    if(stretch) {
        fsetpos(image->fp, &image->fpos );
        fread(image->buf,sizeof(unsigned char),XDIMR,image->fp);
        fread(image->buf,sizeof(unsigned char),XDIMR,image->fp);
    }
}

```

-42-

```

for(i=2;i<(YDIM-2);i++) {
    fread(image->buf,sizeof(unsigned char),XDIMR,image->fp);
    pimage = &image->buf[3];
    for(j=3;j<(XDIM-3);j++) {
        if( *pimage > highest )highest = *pimage;
        if( *pimage < lowest )lowest = *pimage;
        pimage++;
    }
}
if(highest == lowest ){
    printf("something wrong in contrast stretch, zero
contrast");
    exit(1);
}
scale = 255.0 / ( (double)highest - (double)lowest );
dlowest = (double)lowest;
fsetpos(image->fp, &image->fpos );
for(i=0;i<YDIM;i++){
    fgetpos(image->fp, &tmp_fpos );
    fread(image->buf,sizeof(unsigned char),XDIMR,image->fp);
    pimage = image->buf;
    for(j=0;j<XDIM;j++){
        dtemp = ((double)*pimage - dlowest)*scale;
        if(dtemp < 0.0)*(pimage++) = 0;
        else if(dtemp > 255.0)*(pimage++) = 255;
        else *(pimage++) = (unsigned char)dtemp;
    }
    fsetpos(image->fp, &tmp_fpos );
    fwrite(image->buf,sizeof(unsigned
char),XDIMR,image->fp);
}
}

fsetpos(image->fp, &image->fpos );
fsetpos(cortex.fp, &cortex.fpos );
for(i=0;i<YDIM;i++){
    fread(image->buf,sizeof(unsigned char),XDIMR,image->fp);
    fwrite(image->buf,sizeof(unsigned char),XDIMR,cortex.fp);
}
return(1);
}

int clear_int_image(struct int_buf *wordbuffer){
long i,j;
int *pword;
fpos_t tmp_fpos;

fsetpos(wordbuffer->fp, &wordbuffer->fpos );
for(i=0;i<YDIM;i++){
    fgetpos(wordbuffer->fp, &tmp_fpos );
    pword = wordbuffer->buf;
    fread(wordbuffer->buf,sizeof(int),XDIMR,wordbuffer->fp);
    for(j=0;j<XDIM;j++) *(pword++) = 0;
    fsetpos(wordbuffer->fp, &tmp_fpos );
    fwrite(wordbuffer->buf,sizeof(int),XDIMR,wordbuffer->fp);
}
return(1);
}

double find_mean_int(struct int_buf *wordbuffer){
long i,j;
int *pword;
double mean=0.0;

fsetpos(wordbuffer->fp, &wordbuffer->fpos );

```

-43-

```
for(i=0;i<YDIM;i++) {
    pword = wordbuffer->buf;
    fread(wordbuffer->buf,sizeof(int),XDIMR,wordbuffer->fp);
    for(j=0;j<XDIM;j++) mean += (double) *(pword++);
}
mean /= ((double)XDIM * (double)YDIM);

return(mean);
}

int add_uchar_to_int(struct uchar_buf *image,struct int_buf *word){
    unsigned char *pimage;
    int *pword;
    long i,j;
    fpos_t tmp_fpos;

    fsetpos(image->fp, &image->fpos );
    fsetpos(word->fp, &word->fpos );
    for(i=0;i<YDIM;i++){
        pword = word->buf;
        fgetpos(word->fp, &tmp_fpos );
        fread(word->buf,sizeof(int),XDIMR,word->fp);
        pimage = image->buf;
        fread(image->buf,sizeof(unsigned char),XDIMR,image->fp);
        for(j=0;j<XDIM;j++) *(pword++) += (int)*(pimage++);
        fsetpos(word->fp, &tmp_fpos );
        fwrite(word->buf,sizeof(int),XDIMR,word->fp);
    }
    return(1);
}

int add_char_to_uchar_creating_uchar(struct char_buf *cimage,
    struct uchar_buf *image,
    struct uchar_buf *out_image){
    unsigned char *pimage,*pout_image;
    char *pcimage;
    int temp;
    long i,j;

    fsetpos(image->fp, &image->fpos );
    fsetpos(out_image->fp, &out_image->fpos );
    fsetpos(cimage->fp, &cimage->fpos );
    for(i=0;i<YDIM;i++){
        pcimage = cimage->buf;
        fread(cimage->buf,sizeof(char),XDIMR,cimage->fp);
        pimage = image->buf;
        fread(image->buf,sizeof(unsigned char),XDIMR,image->fp);
        pout_image = out_image->buf;
        for(j=0;j<XDIM;j++){
            temp = (int)*(pimage++) + (int)*(pcimage++);
            if(temp<0)temp = 0;
            else if(temp > 255)temp = 255;
            *(pout_image++) = (unsigned char)temp;
        }
        fwrite(out_image->buf,sizeof(unsigned
char),XDIMR,out_image->fp);
    }
    return(1);
}

int copy_int_to_int(struct int_buf *word2,struct int_buf *word){
    long i;

    fsetpos(word2->fp, &word2->fpos );

```

-44-

```

fsetpos(word->fp, &word->fpos );
for(i=0;i<YDIM;i++){
    fread(word->buf,sizeof(int),XDIMR,word->fp);
    fwrite(word->buf,sizeof(int),XDIMR,word2->fp);
}
return(1);
}

void get_snow_images(void){
    unsigned char *psnow,*ptemp;
    int number_snow_inputs;
    int temp,*pword,*pword2,bit;
    long i, j;
    double rms,dtemp;

    live_video(); /* device specific */

    printf("\n\nPlease point camera at a medium lit blank wall. ");
    printf("\nDefocus the lens a bit as well ");
    printf("\nIf possible, place the camera into its highest gain,
and ");
    printf("\nput the gamma to 1.0.");
    printf(" Ensure that the video is not saturated ");
    printf("\nPress any key when ready... ");

    while( !kbhit() );
    printf("\nNow finding difference frame rms value... ");

    /* subtract one image from another, find the rms difference */
    live_video();
    wait_vertical_blanks(2);
    grab_frame(&temp_image);
    live_video();
    wait_vertical_blanks(2);
    grab_frame(&distributed_image); /* use first image as buffer */

    rms = 0.0;
    fsetpos(temp_image.fp, &temp_image.fpos );
    fsetpos(distributed_image.fp, &distributed_image.fpos );
    for(i=0;i<YDIM;i++){
        ptemp = temp_image.buf;
        fread(temp_image.buf,sizeof(unsigned
char),XDIMR,temp_image.fp);
        psnow = distributed_image.buf;
        fread(distributed_image.buf,sizeof(unsigned
char),XDIMR,distributed_image.fp);
        for(j=0;j<XDIM;j++){
            temp = (int) *(psnow++) - (int) *(ptemp++);
            dtemp = (double)temp;
            dtemp *= dtemp;
            rms += dtemp;
        }
    }
    rms /= ( (double)XDIM * (double)YDIM );
    rms = sqrt(rms);
    printf("\n\nAn rms frame difference noise value of %lf was
found.",rms);
    printf("\nWe want at least %lf for good measure",RMS_VAL);
    /* we want rms to be at least RMS_VAL DN, so ... */7
    if(rms > RMS_VAL) number_snow_inputs = 1;
    else {
        dtemp = RMS_VAL / rms;
        dtemp *= dtemp;
        number_snow_inputs = 1 + (int)dtemp;
    }
    printf("\n%d images will achieve this noise
level",number_snow_inputs);
}

```

-45-

```

/* now create each snowy image */
printf("\nStarting to create snow pictures... \n");
fsetpos(snow_images.fp, &snow_images.fpos ); /* set on first
image*/
for(bit = 0; bit < BITS; bit++){
    clear_int_image(&temp_wordbuffer);
    for(i=0;i<number_snow_inputs;i++){
        live_video();
        wait_vertical_blanks(2);
        grab_frame(&temp_image);
        add_uchar_to_int(&temp_image,&temp_wordbuffer);
    }
    clear_int_image(&temp_wordbuffer2);
    for(i=0;i<number_snow_inputs;i++){
        live_video();
        wait_vertical_blanks(2);
        grab_frame(&temp_image);
        add_uchar_to_int(&temp_image,&temp_wordbuffer2);
    }
    /* now load snow_images[bit] with the difference frame
biased by
128 in an unsigned char form just to keep things clean */
    /* display it on cortex also */
    fsetpos(temp_wordbuffer2.fp, &temp_wordbuffer2.fpos );
    fsetpos(temp_wordbuffer.fp, &temp_wordbuffer.fpos );
    fsetpos(temp_image.fp, &temp_image.fpos );
    for(i=0;i<YDIM;i++){
        pword = temp_wordbuffer.buf;
        fread(temp_wordbuffer.buf,sizeof(int),XDIMR,temp_wordbuffer.fp);
        pword2 = temp_wordbuffer2.buf;
        fread(temp_wordbuffer2.buf,sizeof(int),XDIMR,temp_wordbuffer2.fp);
        psnow = snow_images.buf;
        ptemp = temp_image.buf;
        for(j=0;j<XDIM;j++) {
            *(psnow++) = *(ptemp++) = (unsigned char)
(*(pword++) - *(pword2++) + 128);
        }
        fwrite(snow_images.buf,sizeof(unsigned
char),XDIMR,snow_images.fp);
        fwrite(temp_image.buf,sizeof(unsigned
char),XDIMR,temp_image.fp);
    }
    freeze_frame();
    display_uchar(&temp_image,0); /*1 signifies to stretch the
contrast*/
    printf("\rDone snowy %d    ",bit);
    wait_vertical_blanks(30);
}
return;
}

void loop_visual(void){
    unsigned char *psnow;
    char *pcomp;
    long i,j,count = 0;
    int ok=0,temp,bit,add_it;

```

-46-

```

double scale = 1.0 / RMS_VAL;
double dtemp,tmpscale;
fpos_t tmp_fpos;

/* initial rms of each snowy image should be around 5 to 10 DN.
let's assume it is 5, and assume further that our acceptable
noise level of
the full snowy composite should be approximately 1 DN, thus we
need to
scale them down by approximately 5*BITS as a first guess, then
do the
visual loop to zoom in on final acceptable value */

printf("\n\n Now calculating initial guess at amplitude...
\n");
while( !ok ){
    /* calculate snow_composite */
    /* clear composite */
    clear_char_image(&snow_composite);

    fsetpos(snow_images.fp, &snow_images.fpos ); /* set on
first image*/
    for(bit=0;bit<BITS;bit++){
        j = 128 >> bit;
        if( our_code & j)add_it=1;
        else add_it=0;
        fsetpos(snow_composite.fp, &snow_composite.fpos );
        for(i=0;i<YDIM;i++){
            psnow = snow_images.buf;
            fread(snow_images.buf,sizeof(unsigned
char),XDIMR,snow_images.fp);
            fgetpos(snow_composite.fp, &tmp_fpos );

            fread(snow_composite.buf,sizeof(char),XDIMR,snow_composite.fp);
            pcomp = snow_composite.buf;
            for(j=0;j<XDIM;j++) {
                dtemp = ((double)*(psnow++)-128.0) * scale;
                if(dtemp<0.0){
                    temp = -(int) fabs( -dtemp +0.5);
                }
                else {
                    temp = (int) fabs( dtemp +0.5);
                }
                if(temp > 127) {
                    temp = 127;
                }
                else if(temp < -128) {
                    temp = -128;
                }
                if(add_it){
                    *(pcomp++) += (char)temp;
                }
                else {
                    *(pcomp++) -= (char)temp;
                }
            }
            fsetpos(snow_composite.fp, &tmp_fpos );
        }
        fwrite(snow_composite.buf,sizeof(char),XDIMR,snow_composite.fp);
        }
        printf("\rDone snowy %d    ",bit);
    }

/* add snow composite to test image to form dist image */
add_char_to_uchar_creating_uchar(
    &snow_composite,
    &test_image,

```

```

        &distributed_image);

/* display both and cue for putting scale down, up or ok */
i=count = 0;
printf("\n Depress any key to toggle, enter to move on...\n");
");
printf("\r Distributed Image...      ");
display_uchar(&distributed_image,0);
while( getch() != '\r' ){
    if( (count++) % 2){
        printf("\r Distributed Image...      ");
        display_uchar(&distributed_image,0);
    }
    else {
        printf("\r Original Image...      ");
        display_uchar(&test_image,0);
    }
}
printf("\nScale = %lf ",scale);
printf("\nEnter new scale, or >1e6 for ok... ");
scanf("%lf",&tmpscale);
if(tmpscale > 1e6)ok=1;
else scale = tmpscale;
}
/* distributed image now is ok; calculate actual snow_images
used and
store in those arrays; */
fsetpos(snow_images.fp, &snow_images.fpos ); /* set on first
image*/
printf("\nNow storing snow images as used... \n");
for(bit=0;bit<BITS;bit++){
    for(i=0;i<YDIM;i++){
        psnow = snow_images.buf;
        fgetpos(snow_images.fp, &tmp_fpos );
        fread(snow_images.buf,sizeof(unsigned
char),XDIMR,snow_images.fp);
        for(j=0;j<XDIM;j++) {
            dtemp = ((double)*psnow -128.0) * scale;
            if(dtemp<0.0){
                temp = -(int) fabs( -dtemp +0.5);
            }
            else {
                temp = (int) fabs( dtemp +0.5);
            }
            *(psnow++) = (unsigned char)(temp + 128);
        }
        fsetpos(snow_images.fp, &tmp_fpos );
        fwrite(snow_images.buf,sizeof(unsigned
char),XDIMR,snow_images.fp);
    }
    printf("\rDone snowy %d   ",bit);
}
return;
}

}

```

```

double find_grad(struct int_buf *image,int load_buffer2){
    int bufl[XDIMR],buf2[XDIMR],buf3[XDIMR];
    int *pbuf1,*pbuf2,*pbuf3,*p2;
    double total=0.0,dtemp;
    long i, j;
    fpos_t tmp_pos;

```

-48-

```

fsetpos(image->fp, &image->fpos );
fgetpos(image->fp, &tmp_pos );

fsetpos(temp_wordbuffer2.fp, &temp_wordbuffer2.fpos );
for(i=1;i<(YDIM-1);i++){
    fsetpos(image->fp, &tmp_pos );
    fread(buf1,sizeof(int),XDIMR,image->fp);
    fgetpos(image->fp, &tmp_pos );
    fread(buf2,sizeof(int),XDIMR,image->fp);
    fread(buf3,sizeof(int),XDIMR,image->fp);
    pbuf1=buf1;
    pbuf2=buf2;
    pbuf3=buf3;
    p2 = temp_wordbuffer2.buf;

    if(load_buffer2){
        for(j=1;j<(XDIM-1);j++){
            dtemp = (double)*(pbuf1++);
            dtemp += (double)*(pbuf1++);
            dtemp += (double)*(pbuf1--);
            dtemp += (double)*(pbuf2++);
            dtemp -= (8.0 * (double) *(pbuf2++));
            dtemp += (double)*(pbuf2--);
            dtemp += (double)*(pbuf3++);
            dtemp += (double)*(pbuf3++);
            dtemp += (double)*(pbuf3--);
            *p2 = (int)dtemp;
            if( *p2 > GRAD_THRESHOLD ){
                *(p2++) -= GRAD_THRESHOLD;
            }
            else if( *p2 < -GRAD_THRESHOLD ){
                *(p2++) += GRAD_THRESHOLD;
            }
            else {
                *(p2++) = 0;
            }
        }

        fwrite(temp_wordbuffer2.buf,sizeof(int),XDIMR,temp_wordbuffer2.fp);
    }
    else {

        fread(temp_wordbuffer2.buf,sizeof(int),XDIMR,temp_wordbuffer2.fp);
        for(j=1;j<(XDIM-1);j++){
            dtemp = (double)*(pbuf1++);
            dtemp += (double)*(pbuf1++);
            dtemp += (double)*(pbuf1--);
            dtemp += (double)*(pbuf2++);
            dtemp -= (8.0 * (double) *(pbuf2++));
            dtemp += (double)*(pbuf2--);
            dtemp += (double)*(pbuf3++);
            dtemp += (double)*(pbuf3++);
            dtemp += (double)*(pbuf3--);

            dtemp -= (double) *(p2++);
            dtemp *= dtemp;
            total += dtemp;
        }
    }
}

return(total);
}

```

-49-

```

void search_1(struct uchar_buf *suspect){
    unsigned char *psuspect,*psnow;
    int bit,*pword,temp;
    long i,j;
    double add_metric,subtract_metric;
    fpos_t tmp_fpos;

    /* this algorithm is conceptually the simplest. The idea is to
step      through each bit at a time and merely see if adding or
subtracting the
individual snowy picture minimizes some 'contrast' metric.
This should be the most crude and inefficient, no where to go
but      better */

    fsetpos(snow_images.fp, &snow_images.fpos );
    temp=256;
    clear_int_image(&temp_wordbuffer);
    add_uchar_to_int(suspect,&temp_wordbuffer);
    find_grad(&temp_wordbuffer,1); /* 1 means load temp_wordbuffer2
*/
    for(bit=0;bit<BITS;bit++){
        /* add first */
        fgetpos(snow_images.fp, &tmp_fpos );
        fsetpos(suspect->fp, &suspect->fpos );
        fsetpos(temp_wordbuffer.fp, &temp_wordbuffer.fpos );
        for(i=0;i<YDIM;i++){
            pword = temp_wordbuffer.buf;
            psuspect = suspect->buf;
            psnow = snow_images.buf;
            fread(suspect->buf,sizeof(unsigned
char),XDIMR,suspect->fp);
            fread(snow_images.buf,sizeof(unsigned
char),XDIMR,snow_images.fp);
            for(j=0;j<XDIM;j++){
                *(pword++)=(int)*(psuspect++)+(int)*(psnow++)-128;
            }
            fwrite(temp_wordbuffer.buf,sizeof(int),XDIMR,temp_wordbu
ffer.fp);
        }
        add_metric = find_grad(&temp_wordbuffer,0);

        /* then subtract */
        fsetpos(snow_images.fp, &tmp_fpos );
        fsetpos(suspect->fp, &suspect->fpos );
        fsetpos(temp_wordbuffer.fp, &temp_wordbuffer.fpos );
        for(i=0;i<YDIM;i++){
            pword = temp_wordbuffer.buf;
            psuspect = suspect->buf;
            psnow = snow_images.buf;
            fread(suspect->buf,sizeof(unsigned
char),XDIMR,suspect->fp);
            fread(snow_images.buf,sizeof(unsigned
char),XDIMR,snow_images.fp);
            for(j=0;j<XDIM;j++){
                *(pword++)=(int)*(psuspect++)-(int)*(psnow++)+128;
            }
            fwrite(temp_wordbuffer.buf,sizeof(int),XDIMR,temp_wordbu
ffer.fp);
        }
        subtract_metric = find_grad(&temp_wordbuffer,0);

        printf("\nbit place %d: add=%le ,
sub=%le",bit,add_metric,subtract_metric);
        temp/=2;
        if(add_metric < subtract_metric){
            printf(" bit value = 0");
        }
    }
}

```

-50-

```

        }
    else {
        printf(" bit value = 1");
        found_code += temp;
    }
printf("\n\nYour magic number was %d",found_code);
return;
}

void search_2(unsigned char *suspect){
if(suspect);
return;
}

void loop_simulation(void){
unsigned char *ptemp,*pdist,
int *pword,int_mean,ok=0,temp;
long i,j;
double mean,scale;

/* grab a noisy image into one of the temp buffers */
printf("\ngrabbing noisy frame...\n");
clear_int_image(&temp_wordbuffer);
for(i=0;i<NUM_NOISY;i++){
    live_video();
    wait_vertical_blanks(2);
    grab_frame(&temp_image);
    add uchar_to_int(&temp_image,&temp_wordbuffer);
    j=(long)NUM_NOISY;
    printf("\r%ld of %ld      ",i+1,j);

    /* find mean value of temp_wordbuffer */
    mean = find_mean_int(&temp_wordbuffer);
    int_mean = (int)mean;

    /* now we will add scaled version of this 'corruption' to our
distributed
image */
    scale = 1.0;
    while( !ok ){
        /* add noise to dist image storing in temp_image */
        fsetpos(distributed_image.fp, &distributed_image.fpos );
        fsetpos(temp_wordbuffer.fp, &temp_wordbuffer.fpos );
        fsetpos(temp_image.fp, &temp_image.fpos );
        for(i=0;i<YDIM;i++){
            pdist = distributed_image.buf;
            pword = temp_wordbuffer.buf;
            ptemp = temp_image.buf;
            fread(distributed_image.buf,sizeof(unsigned
char),XDIMR,distributed_image.fp);
            fread(temp_wordbuffer.buf,sizeof(int),XDIMR,temp_wordbuf
fer.fp);
            for(j=0;j<XDIM;j++){
                temp = (int) *(pdist++) + *(pword++) - int_mean;
                if(temp<0)temp = 0;
                else if(temp > 255)temp = 255;
                *(ptemp++) = (unsigned char)temp;
            }
        }
    }
}

```

-51-

```

        fwrite(temp_image.buf,sizeof(unsigned
char),XDIMR,temp_image.fp);
    }

    /* display the dist image and the corrupted image */
    display_uchar(&temp_image,0);

    /* apply new 'corrupted' image to search algorithm 1 for id
value */
    search_1(&temp_image);

    /* apply new 'corrupted' image to search algorithm 2 for id
value */
    /*
    search_2(temp_image);
    */

    /* prompt for upping noise content or ok */
    ok = 1;
}

return;
}

int initialize_everything(void) {
    long i,j;
    unsigned char *pucbuf;
    char *pcbuf;
    int *pibuf;

    /* initialize cortex */
    strcpy(cortex.filename,"f:image");
    if((cortex.fp=fopen(cortex.filename,"rb"))==NULL) {
        system("v f g");
    }
    else fclose(cortex.fp);
    if( (_inp(PORT_BASE) == 0xFF) ){
        printf("oops ");
        exit(0);
    }
    /* open cortex for read and write */
    if((cortex.fp=fopen(cortex.filename,"rb+"))==NULL) {
        printf(" No good on open file joe ");
        exit(0);
    }
    fgetpos(cortex.fp, &cortex.fpos );

    /* test_image; original image */
    strcpy(test_image.filename,"e:tst_img");
    if((test_image.fp=fopen(test_image.filename,"wb"))==NULL) {
        printf(" No good on open file joe ");
        exit(0);
    }
    pucbuf = test_image.buf;
    for(i=0;i<XDIM;i++) *(pucbuf++)=0;
    for(i=0;i<YDIM;i++) {
        fwrite(test_image.buf,sizeof(unsigned
char),XDIMR,test_image.fp);
    }
    fclose(test_image.fp);
    if((test_image.fp=fopen(test_image.filename,"rb+"))==NULL) {
        printf(" No good on open file joe ");
        exit(0);
    }
}

```

-52-

```
        }
        fgetpos(test_image.fp, &test_image.fpos );
/* snow_composite; ultimate image added to original image */
strcpy(snow_composite.filename,"e:snw_cmp");
if((snow_composite.fp=fopen(snow_composite.filename,"wb"))==NULL){
    printf(" No good on open file joe ");
    exit(0);
}
pcbuf = snow_composite.buf;
for(i=0;i<XDIM;i++)*(pcbuf++)=0;
for(i=0;i<YDIM;i++){
    fwrite(snow_composite.buf,sizeof(char),XDIMR,snow_composite.fp);
}
fclose(snow_composite.fp);

if((snow_composite.fp=fopen(snow_composite.filename,"rb+"))==NULL){
    printf(" No good on open file joe ");
    exit(0);
}
fgetpos(snow_composite.fp, &snow_composite.fpos );
/* distributed_image; test_img plus snow_composite */
strcpy(distributed_image.filename,"e:dst_img");
if((distributed_image.fp=fopen(distributed_image.filename,"wb"))==NU
LL){
    printf(" No good on open file joe ");
    exit(0);
}
pucbuf = distributed_image.buf;
for(i=0;i<XDIM;i++)*(pucbuf++)=0;
for(i=0;i<YDIM;i++){
    fwrite(distributed_image.buf,sizeof(unsigned
char),XDIMR,distributed_image.fp);
}
fclose(distributed_image.fp);

if((distributed_image.fp=fopen(distributed_image.filename,"rb+"))==N
ULL){
    printf(" No good on open file joe ");
    exit(0);
}
fgetpos(distributed_image.fp, &distributed_image.fpos );

/* temp_image; buffer if needed */
strcpy(temp_image.filename,"e:temp_img");
if((temp_image.fp=fopen(temp_image.filename,"wb"))==NULL){
    printf(" No good on open file joe ");
    exit(0);
}
pucbuf = temp_image.buf;
for(i=0;i<XDIM;i++)*(pucbuf++)=0;
for(i=0;i<YDIM;i++){
    fwrite(temp_image.buf,sizeof(unsigned
char),XDIMR,temp_image.fp);
}
fclose(temp_image.fp);
if((temp_image.fp=fopen(temp_image.filename,"rb+"))==NULL){
    printf(" No good on open file joe ");
    exit(0);
}
fgetpos(temp_image.fp, &temp_image.fpos );
/* temp_wordbuffer; 16 bit image buffer for averaging */
```

-53-

```

strcpy(temp_wordbuffer.filename, "e:temp_wrd");
if((temp_wordbuffer.fp=fopen(temp_wordbuffer.filename,"wb"))==NULL) {
    printf(" No good on open file joe ");
    exit(0);
}
pibuf = temp_wordbuffer.buf;
for(i=0;i<XDIM;i++)*(pibuf++)=0;
for(i=0;i<YDIM;i++) {
    fwrite(temp_wordbuffer.buf,sizeof(int),XDIMR,temp_wordbuffer.fp);
}
fclose(temp_wordbuffer.fp);

if((temp_wordbuffer.fp=fopen(temp_wordbuffer.filename,"rb+"))==NULL)
{
    printf(" No good on open file joe ");
    exit(0);
}
fgetpos(temp_wordbuffer.fp, &temp_wordbuffer.fpos );
/* temp_wordbuffer2; /* 16 bit image buffer for averaging */
strcpy(temp_wordbuffer2.filename,"e:tmp_wrd2");
if((temp_wordbuffer2.fp=fopen(temp_wordbuffer2.filename,"wb"))==NULL)
{
    printf(" No good on open file joe ");
    exit(0);
}
pibuf = temp_wordbuffer2.buf;
for(i=0;i<XDIM;i++)*(pibuf++)=0;
for(i=0;i<YDIM;i++) {
    fwrite(temp_wordbuffer2.buf,sizeof(int),XDIMR,temp_wordbuffer2.fp);
}
fclose(temp_wordbuffer2.fp);

if((temp_wordbuffer2.fp=fopen(temp_wordbuffer2.filename,"rb+"))==NULL)
{
    printf(" No good on open file joe ");
    exit(0);
}
fgetpos(temp_wordbuffer2.fp, &temp_wordbuffer2.fpos );
/* snow_images; BITS number of constituent snowy pictures */
strcpy(snow_images.filename,"snw_imgs");
if((snow_images.fp=fopen(snow_images.filename,"wb"))==NULL) {
    printf(" No good on open file joe ");
    exit(0);
}
pucbuf = snow_images.buf;
for(i=0;i<XDIM;i++)*(pucbuf++)=0;
for(j=0;j<BITS;j++) {
    for(i=0;i<YDIM;i++) {
        fwrite(snow_images.buf,sizeof(unsigned
char),XDIMR,snow_images.fp);
    }
}
fclose(snow_images.fp);
if((snow_images.fp=fopen(snow_images.filename,"rb+"))==NULL) {
    printf(" No good on open file joe ");
    exit(0);
}
fgetpos(snow_images.fp, &snow_images.fpos );

return(1);
}

```

-54-

```
int close_everything(void){
    fclose(test_image.fp);
    fclose(snow_composite.fp);
    fclose(distributed_image.fp);
    fclose(temp_image.fp);
    fclose(temp_wordbuffer.fp);
    fclose(temp_wordbuffer2.fp);
    fclose(snow_images.fp);

    return(1);
}

main(){
    int i,j;

    printf("\nInitializing...\n\n");
    initialize_everything(); /* device specific and global mallocs
*/
    live_video();

    /* prompt for which of the three demos to perform */
    while( demo < 1 || demo > NUM_DEMOS){
        printf("Which demo do you want to run?\n\n");
        printf("1: Digital Imagery and Very High End Photography
Simulation\n");
        printf("2: Pre-exposed Print Paper and other Duplicating\n");
        printf("3: Pre-exposed Original Film (i.e. In-Camera)\n");
        printf("\nEnter number and return: ");
        scanf("%d",&demo);
        if(demo < 1 || demo > NUM_DEMOS){
            printf("\n eh eh ");
        }
    }

    /* acquire test image */
    printf("\nPress any key after your test scene is ready... ");
    getch();
    grab_frame(&test_image); /*grab_frame takes care of device
specific stuff*/

    /* prompt for id number, 0 through 255 */
    printf("\nEnter any number between 0 and 255.\n");
    printf("This will be the unique magic code placed into the
image: ");
    scanf("%d",&our_code);
    while(our_code<0 || our_code>256){
        printf(" Between 0 and 255 please ");
        scanf("%d",&our_code);
    }

    /* feed back the binary code which will be embedded in the image
*/
    printf("\nThe binary sequence ");
    for(i=0;i<BITS;i++){
        j = 128 >> i;
        if( our_code & j)printf("1");
        else printf("0");
    }
    printf(" (%d) will be embedded on the image\n",our_code);
    /* now generate the individual snow images */
    get_snow_images();
}
```

-55-

```
loop_visual(); /* this gives visual feedback on 'tolerable'
noise level */
printf("\nWe're now to the simulated suspect... \n");
loop_simulation();
closeEverything();
return(0);
}
```

Claims

1. A method of identification coding an input signal so as to permit its later identification, the method including the steps:
 - modulating a noise signal with a code number to produce a signature signal; and modulating the input signal with the signature signal to produce an identification coded output signal;
 - wherein the coded output signal can be analyzed to discern the code number with which it was modulated.
2. A method of identification coding an input signal so as to produce an encoded output signal, the input signal being a quantized signal having inherent noise, said signal corresponding to aural or visual information, the identification coding of the output signal preserving the corresponding aural/visual information without human-perceptible degradation, the identification coding permitting later identification of the output signal, the method including modulating a noise signal with a code number to produce a signature signal, adding the signature signal to the input signal to produce an identification coded output signal, the signature signal having an amplitude below a threshold of human aural/visual perceptibility when added to the input signal, the adding step effecting distribution of the signature signal throughout the entirety of the output signal.
3. A method of data processing including: providing a digital carrier signal, and modulating the digital carrier signal to imperceptibly embed an identification signal thereon, the method characterized by: compressing the modulated digital carrier signal with lossy data compression to produce a compressed signal, decompressing the compressed signal, and discerning the embedded identification signal from the decompressed signal, wherein the lossy data compression does not preclude recovery of the embedded identification signal.
4. An apparatus for encoding a sampled input signal, the sampled input signal having inherent noise, the apparatus including an input terminal, a digital noise source, storage for an identification code word, means for maintaining a pointer to a bit of the identification code word, an adder, and an output terminal, the input terminal being coupled to a first input of the adder, the noise source being coupled to a second input of the adder, the pointer providing said bit of the identification code word to a control input of the adder, an output of the adder being coupled to the output terminal.
5. The apparatus of claim 4 which further includes a look-up table, a first scaler, a second scaler, a scale control device, and a memory, the look-up table having an input coupled to the input terminal, one of said scalers having a control input coupled to an output of the look-up table, the other of said scalers having a control input coupled to the scale control device, said scalers being serially interposed between the noise source and the adder, the memory having an input coupled to a location between the noise source and the second input of the adder.

6. A method of identification coding a sampled input signal, the sampled input signal having inherent noise, the method comprising:

providing an N-bit code number;

for each of a plurality of samples of the input signal:

(a) providing a sample of a time- or spatially-varying modulation signal;

(b) selecting one bit of the N-bit code number; and

(c) if said bit has a first value, adding the modulation signal sample to the sample of the input signal, yielding a sample of an identification coded output signal.

7. The method of claim 6 which includes performing steps (a) - (c) for each sample of the input signal.

8. The method of claim 6 which further includes storing, for later use, data from which the modulation signal sample can be reconstructed.

9. The method of claim 6 which includes generating the time-varying modulation signal sample by providing a pseudo-random number and weighting said number with a scaling factor, said scaling factor being a function of the input signal sample.

10. The method of claim 6 which includes selecting the one bit of the N-bit code number by cycling through the number, advancing one bit position for each successive sample of the input signal.

11. The method of claim 6 which further includes:

if said selected bit of the N-bit code number has a second value, subtracting the modulation signal sample from the sample of the input signal, yielding a sample of the identification coded output signal.

12. Storage medium having stored thereon a signal processed in accordance with the method of claim 6.

13. The invention of claim 12 in which the storage medium is a magnetic medium.

14. The invention of claim 12 in which the storage medium is a printed medium.

15. The invention of claim 12 in which the storage medium is a compact disk (CD).

16. A method of identification coding each of a plurality of samples of a sampled input signal, the input signal having inherent noise, characterized by:
using the sample of the input signal to obtain a scaling factor uniquely associated thereto;

weighting a signature datum in accordance with said scaling factor; and
modulating the sample of the input signal in accordance with said weighted signature datum.

17. The method of claim 16 in which the scaling factors increase monotonically with the values of the input signal samples with which they are associated.

18. The method of claim 16 in which a four-fold increase in a value of the sampled input signal corresponds to approximately a two-fold increase in a value of the scaling factor associated therewith.

19. A method of processing a sampled input signal with an N-bit signature word to produce an identification-coded output signal, the sampled input signal having inherent noise, wherein the complete N-bit signature finds expression M times in an excerpt of the identification-coded output signal having a length of M^*N samples, for some value of M greater than one.

20. The method of claim 19 characterized by processing each sample of the input signal in accordance with at least part of the signature word.

21. In a method of processing a source signal that includes a number of elements, each with an associated value, an improvement characterized by altering the source signal in accordance with an embedded signal so as to encode an identification code therein, the embedded and altered signals each including a number of elements, each with an associated value, wherein an element of the altered signal has a value different than that of corresponding elements in both the source and embedded signals, and in which the identification code and certain pseudo-random reference data are used to generate the embedded signal, the association between the embedded signal and the identification code being undiscernible without availability of the reference data.

22. In a method of processing a source signal that includes a number of elements, each with an associated value, an improvement characterized by:

providing an N bit digital identification code, each bit having a "1" or "0" value;

providing N different reference signals, one being associated with each bit position in the digital identification code;

summing the reference signals for which the corresponding bit position in the identification code has a "1" value, thereby producing an embedded signal;

altering the source signal in accordance with the embedded signal so as to encode an identification code therein;

the embedded and altered signals each including a number of elements, each with an associated value, wherein an element of the altered signal has a value different than that of corresponding elements in both the source and embedded signals.

FIG. 4






FIG. 1



2 / 7

FIG. 2



3 / 7

FIG. 3




FIG. 5





FIG. 6







FIG. 9A




FIG. 9B




FIG. 9C

BORDER
CONTINUITY
404




FIG. 10




FIG. 11





FIG. 12



THIS PAGE BLANK (USPTO)